

■ HOTTEST NEWS ■ LATEST REVIEWS ■ BEST TUTORIALS ■ YOUR PROBLEMS SOLVED

LINUX

FORMAT

THE UK'S BEST-SELLING LINUX MAGAZINE!



DO MORE!

WE SHOW YOU HOW WITH OVER 40 PAGES OF TUTORIALS:



■ Installing software ■ Text in GIMP ■ Samba filesharing ■ PHP
■ Realsoft3D ■ MP3 ■ KDE ■ Blender AND MORE! **p59**

REVIEWED INSIDE:



■ Neoware
thin client
appliance



■ JBuilder 9
Java dev
studio



■ SuSE
Enterprise
Desktop



■ Ximian
Desktop
platform



■ Navaho
integrated
server

...AND MORE PRODUCT REVIEWS STARTING ON **p17**

CORPORATE DESKTOP

SuSE distro sets out to win hearts and minds in the office **p24**

PROBLEM SOLVED!

Seven pages of the best advice – our experts have the answers **p94**

ACHIEVE BETTER SOUND



Get the most from your audio hardware – extensive guide to sound servers, streaming and more **p45**

LATEST REVIEWS:

Neoware, JBuilder9, SuSE Enterprise Desktop & more **p17**

“RSYNC WILL SAVE YOU TIME, EFFORT AND EXPENSIVE MISTAKES WHEN IT COMES TO BACKUPS” p52

DVD issue also available Printed in the UK
LXF44 SEPTEMBER 2003 £5.99

 9 771470 423019 09
 (£5.99 outside UK and Republic of Ireland)

www.linuxformat.co.uk

You want to do what?

Motto-wise, the average Linux user seems to live by the maxim "if it can be done, it can be done better on Linux". For a lot of uses, this probably turns out to be true. And what's also great about using Linux is that even if you don't quite succeed in what you are trying to accomplish, it's rarely a wasted experience, because of the knowledge you pick up along the way. Sometimes though, the quest becomes frustrating – what you are trying to do just doesn't seem to have been done before, or if it has, nobody has thought to leave behind a HOWTO. It is surprising the number of key parts of Linux that are either under-documented or occasionally not documented at all. Overcoming the odds in these situations may be character-forming, but can also be frustrating when you simply need to get things done. We are constantly addressing this problem, by delivering comprehensive and comprehensible tutorials on virtually anything you might like to do, and as the cover this issue suggests, we've managed to cram a few more in for you this

month! With several new series starting soon, we think we have the key subjects covered (and remember, plenty of past tutorials are available online in the Archive section of our website at www.linuxformat.co.uk). If you have any suggestions for things we haven't covered, or comments on the tutorials we have run, please do drop us a line.

Our big feature this issue also has a practical bent. The upcoming kernel 2.6 is set to rely on ALSA by default, instead of OSS. If you are interested in sound and music, how to get more from your hardware and how to prepare for the new kernel release, you should read this.

And if you are more interested in what other people can do for you, there's a plethora of interesting new products reviewed, from thin clients and server appliances to SuSE's bold new gambit for the corporate desktop. As usual, we go into levels of detail you won't find anywhere else, so if you want the real story, get reading. Hope you enjoy the issue!





Nick Veitch EDITOR

Whatever your Linux skill level, we've got the tutorials you need to do even more **p59**

SuSE Enterprise Desktop 1 helps businesses migrating from Windows **p24**

If you enjoy games or multimedia, why not boost Linux's sound and annoy your neighbours? **p45**



MEET LINUX FORMAT'S TEAM OF WRITERS...



Andrew Channelle
Andy is the LXF chap with a nose for news, and is the novices' best friend with his great Beginners' tutorials.

David Cartwright
Veteran journalist and Linux consultant, he eats and sleeps real-world Linux usage and Open Source advocacy.



David Coulson
Our LXF Answers guy is a networking and security guru with plenty of sysadmin experience.

Hoyt Duff
Author of several *Red Hat Unleashed* books and long-time LXFer, if you can't find him, he's probably gone fishing.



Richard Drummond
Lightning fried his PC, but he still got his Better Sound feature in on time. What a hero!

Mike Saunders
Now considered an LXF veteran, he's the only writer to have caused the editor to laugh out loud. Intentionally...



Jono Bacon
Jono is a core KDE developer, writer, web developer, musician and sound engineer. And insomniac...

Dr Chris Brown
A freelance Linux writer and Unix instructor. He has a PhD in Particle Physics, but hopes it doesn't show.



Paul Hudson
Having pretended to be a novice when calling web hosting companies, he's now even less modest about his hacking skills!

Michael J Hammel
Professional GIMP artist who pens (or pencils) our current Open Source graphics tour-de-force.



AIMS OF THE MAGAZINE

Linux Format is a magazine dedicated to Linux and the Open Source community. We aim:

- To provide the most accurate, unbiased and up to date information on all things Linux.
- To promote the use of Linux in business and the home, for servers and on the desktop.
- To support the Open Source community by providing a resource of information, and a forum for debate.
- To help all readers get more from their Linux experience by providing insightful and useful tutorials.

CONTACT US

Letters for publication:
linuxformat@futurenet.co.uk

Subscriptions/back issues:
subs@futurenet.co.uk

Technical help/Ask the Experts:
linuxformat@futurenet.co.uk

Disc problems:
linuxformat@futurenet.co.uk

General enquiries:
linuxformat@futurenet.co.uk

Website: www.linuxformat.co.uk

Or send your letters to:
LINUX Format, Future Publishing,
30 Monmouth Street, Bath, BA1 2BW
Phone: 01225 442244
Fax: 01225 732295

More contact info on p114

LXF44 September 2003

Contents



Welcome to another jam-packed issue of *Linux Format*, your guide to all things Linux!



COVER HIGHLIGHTS

Do more!

59

Take your Linux skills to the limit with 40 pages of tutorials!

24 SuSE Enterprise Desktop 1

A new distro to boost your office productivity

33 ROUNDUP: Text Editors

The most controversial Linux subject ever?

45 Make Linux sound better

Turning your box all the way up to eleven...

52 What On Earth is... rsync?

Synchronise your files without any 'rsync about'!



'Jargon sets users apart from those not like them... who's part of the group if they pronounce the word "Line-ucks"?'

07 Divisiveness can have its uses...

'If I, having used Linux for a long time, had problems with certain things, then newcomers would have no chance!'

37 The 'best text editor' debate rages on!

LINUXPRO

Who's stolen your domain name? Web address hijacking and what you can and can't do about it **Compiler Writing** Flex and Bison Q&A **Web Hosting** We test and rate the top six UK Linux web providers



29 Looks good and is no mean performer either...



6 News

Analysis of the latest Linux stories

12 Mailserver

Your letters and Linux tips

18 Neoware Eon review

Stealthy thin client hardware

20 Seapine SCM review

A real alternative to CVS

22 Borland JBuilder 9 review

Rapid Application Development

24 SuSE Enterprise Desktop review

Version 1 of a new business distro

26 Ximian Desktop 2 review

A great place for novices to start

29 Navaho TeamCAT server review

A thoroughly impressive bit of kit

30 Book reviews

Java, C#, Python and FreeBSD

38 HotPicks

The latest and greatest software

60 MP3 jukebox tutorial

Build your own music file manager

62 Red Hat package management

Alternatives to the official tools

70 Samba tutorial

Authorisation and authentication

74 Realsoft 3D tutorial

Get started with this great modeller

78 Blender tutorial: Animation

Animation, that's what you need!

82 The GIMP tutorial

Adding text to your pictures

86 PHP tutorial

Optimising your code

90 Beginners' Guide — KDE

The K Desktop Environment

94 Answers

Your queries answered

110 User Groups

Local support and fun


18




38



74



52



COVERDISCS

A DVD or 2 CDs packed full of the latest Linux goodies **104**



CDS A AND B

Multimedia Create, convert and enjoy audio and video; **Ogle, Xine, MPlayer, Video LAN Client** video and DVD players; **LIVES, Kino, Cinelerra, DVD::Rip, Grip, GSMP** editing and conversion apps; **XMMS, TiMidity++, Zinf** audio players; **The GIMP, ImageMagick** graphics editors; **Realsoft 3D** see tutorial page 74; **wxPython** give your Python progs GUIs; **TuxPaint** fun and educational art for kids



DVD

JAMD Linux distro based on Red Hat plays multimedia 'out of the box'; **KOffice 1.3** beta version of this great office suite; **LBreakout2** Retro-style arcade action **Mozilla 1.4, Opera 7.11**

Please read the coverdisc instructions on page 107 before installing from coverdiscs!



SAVE MONEY!

SUBSCRIBE TO LINUX FORMAT

See page 102 or phone 0870 4448645



Newsdesk

● IBM desktide server ● HP Linux enterprise PCs ● SuSE Linux hardware from Wal-Mart ● SCO News
● Scribus Linux DTP ● Linare distro ● Red Hat unboxed! ● Embedded Linux News ● Linux on Xbox

2.6 PRE-X KERNEL OUT IN THE WILD

Linux 2.6 coming soon

Just weeks after announcing his migration from Transmeta to Open Source Development Labs (OSDL), Linus Torvalds has begun the process of releasing test versions of the next iteration of the Linux kernel. The 2.6pre-X replaces the 2.5 development series, signalling a feature freeze and the bug squashing process that, with the last major release, took seven months. "The point of the test versions is to make more people realise that they need testing and get some straggling developers to realise that it's too late to worry about the next big feature," he wrote on the Linux Kernel Mailing List. Torvalds signalled that he

Whether you run big irons or a desktop box, 2.6 will boost Linux performance.



expected the testing phase to be passed in less time: "I hope (and believe) that we have fewer issues facing us in the current 2.6.0."

Maintenance of 2.6 has now been handed over to fellow OSDL staffer Andrew Morton which, according to developer Alan Cox, may speed the release process as Morton is less likely to let new features worm their way into the kernel. "Linus faced with a neat new feature is rather like a small child faced with chocolate," he told InfoWorld.

In an interview with Australia's ComputerWorld, Morton said the 2.5 code is in better shape than previous releases, but that there were still 'many months worth of bugfixing' before the 2.6 series is stable enough to be used in a production environment. As a "mad guess" Morton predicted that an official release was maybe three months away.

Torvalds called on distribution developers to offer the test kernel to users as soon as possible, and upgrade internal systems. Red Hat became the first of the major sellers to offer packaged upgrades of the kernel for

Red Hat 9 (Shrike) users, while SuSE says its version will be available shortly. A SuSE spokesperson said the company expected to have 'enterprise ready' apps built on 2.6 available within 12 months.

Spot the difference

Morton says that, for ordinary users, changes will be subtle. "People do say that the 2.5 kernel is smoother and a little faster in desktop situations," he said, but most improvement will be noticeable for high-end users. "2.6 will be much more scaleable on large SMP machines than 2.4. 16 CPUs is a reasonable target, and we seem to be doing fairly well on 32-processor servers."

High-end features include support for Intel's Physical Address Extension (PAE) for better memory management, efficiency gains in the filesystem, a boost in the number of unique users and groups available from 65,000 to 4 billion and support for Non-Uniform Memory Access (NUMA) servers. This last is seen as so significant that some developers have suggested the new kernel merits a leap to version 3.0.

In his excellent précis of the new kernel's features, Joseph Pranevich says that NUMA "is a step beyond Symmetric MultiProcessing (SMP) in the multi-processing world" offering greater efficiency on systems that have many processors. The most significant gain, he says, is that NUMA doesn't repeat the limitation imposed by traditional uniprocessor and SMP systems, which means a 'single pool of memory is equally accessible from all processors'.

"NUMA servers leap beyond that by introducing the concept that, for a specific processor, some memory is closer than others... In many ways, the new NUMA architecture is an example of a very tight-knit cluster."

While scaling up is seen as the primary goal, effort has been made to ensure the new kernel works well on the latest generation of processors including Motorola's m68k (Dragonball), and AMD's new 64bit Opteron. There are also hardware optimisations for Crusoe, Pentium 4 (incl. better Hyperthreading support), Xeon, Pentium 3-M, and Pentium 4-M, and AMD processors.

Links

Mortals looking for information about the new kernel are well catered for by Joseph Pranevich at:


<http://www.kniggit.net/www26.html>

Coders meanwhile can get more detailed information by Dave Jones at:

<http://www.codemonkey.org.uk/post-halloween-2.5.txt>

Dedicated kernel watchers can get the latest development news from the Linux Kernel Mailing List and other sources at Kernel Traffic from:

<http://kt.zork.net/kernel-traffic/>



MUSE handles up to six audio streams.

BROADCAST NEWS

The no footprint radio show from a CD!

Hot on the heels of the various

Knoppix task-centric live distributions comes DyneBolic Linux, designed to be a fully-fledged broadcast studio capable of running from the CD. The developers say the project is aimed at media activists, artists and other creatives and can be used to create, manipulate and distribute audio and video using only free software. One of the developers even had the system running on Xbox!

DynaBolic has been developed mainly as a vehicle for MUSE, the Multiple Streaming Engine program which can handle up to six discrete audio inputs and stream them across a LAN or the Internet. Its creators Jaromil, Bomboclat and c1cc10, say the live CD is ideal for users wishing to create their own Internet radio stations. "Put in the CD, reboot and start your show!"

ANIMATION

Dreamworks goes 100 per cent Linux

Linux has taken another leap into the forefront of the movie world with the release of *Sinbad: Legend of the Seven Seas*, the first film to be created entirely on Linux systems. After successfully modelling and rendering most of *Shrek*, Dreamworks invested heavily in graphical workstations from Hewlett Packard running Red Hat to marry traditional 2D animation techniques to cutting-edge 3D work.

The biggest challenge, HP's Martin Fink said, was the ocean-bound nature of the story. "Creating a movie that is centred on scenes with water is a huge technical challenge in digital animation," he said. The solution was to forget about rendering the waves on a frame-by-frame basis and, instead, create a fully-fledged digital ocean. Animators could then select sections of the model and add landscape elements. Following the



Creating a digital ocean shaved 'years' off the production time.

traditional method, HP's Mike Balma told *Wired*, would have added years to the film making process.

Ed Leonard, Dreamworks' head of animation technology said the combination of low-cost hardware and the performance gains of Linux meant the film was hitting the limits of current technology. "We are creating movies that would have been impossible to make as little as a year ago," he said.

NEWSBYTES

■ Coventry's Altair Engineering has been granted exclusive rights to provide Linux clusters based on NEC's High Performance Computer range to the manufacturing sector in the UK and Ireland. Altair and NEC have long-running track records, especially in the automotive industry, and the partnership, Altair's Dr Royston Jones said, would allow companies to benefit from the explosion in virtual design which is now possible.

■ Microsoft has moved Linux up to threat number two, with the risk of general economic downturn at the top spot. MS is said to be responding to the challenge on a number of fronts, including lowering the cost of Windows CE and allowing developers not just to view and modify the code, but also offer bespoke products based on it.

■ A recent edition of USA Today, one of the most widely read newspapers in the US, featured a surprisingly in-depth investigation into the awarding of the contract to supply 14,000 for the city of Munich. The article (archived at www.usatoday.com/usatoday/20030714/5320229s.htm) suggests that Microsoft offered to unbundle *Word* from the *MS Office* suite and drop the price to prevent the business going to SuSE and IBM, and may even have gained access to the original Unilog study which set out the needs of the city. Access to this document, denied by Microsoft Germany's spokesman, would have broken the rules of the bid.

■ The British Computer Society's Information Security Specialist Group will be hosting a Linux Day to introduce members to the benefits and risks of using Open Source software. Topics covered at the event include a Linux primer for non-technical users, hardening Linux and the correct use of firewalls. The conference will be held on 25 September 2003 at the Middle East Association, 33 Bury Street St James's, London SW1Y 6AX. Bookings can be made by emailing kateboncz@aol.com.

■ Trolltech has released a new scripting interface for Qt. *Qt Script for Applications (QSA) 1.0* brings full script functionality to any Qt based applications with, Trolltech CEO Eirik Eng says, the addition of a few lines of code "and then the applications are enabled for scripting and extending." The language itself is based on ECMAScript and so should feel familiar to users of other languages such as Javascript or Microsoft's Jscript.

■ If you're not that enamoured with the recent behaviour of the RIAA, you might enjoy reading the info and using the downloadable graphics here: <http://techfocus.org/comments.php?id=3662&catid=17>

Hoyt Duff

The author is one of 800 Hoyts living in the USA and runs a little fishing pier when he's not dabbling with his computers.



COMMENT

It's how you speak

Any person educated in the arts or sciences will tell you that an important part of learning their discipline is learning the language of that discipline. To the rest of the world, these words may sound like gibberish, in that field of endeavour, they have specific meaning; those words allow for clear communication.

That unique vocabulary, also referred to as "jargon", also serves to set the users apart from those not like them and permits those so initiated to identify each other. Who's part of the group if they pronounce the word "Line-uicks"? Not that person, as the preferred pronunciation marks a true member of the group. Likewise, there are other words whose correct usage marks an individual as a part of the group.

The notorious Eric S. Raymond (aka ESR) publishes at <http://catb.org/~esr/jargon/html/index.html> a lexicon known as The Jargon File. It's part history, part opinion, part dictionary much in the spirit of Ambrose Bierce's Devil's Dictionary. It runs the gamut from the hacker/cracker debate to the sexual habits and dress of hackers. The words are contained in the Glossary where you will find such words as baz, biff, feep, lenna, spl, and zorch as well as more common words like flamobait.

Thanks to the DICT protocol, these definitions as well as those from several other dictionaries are available from *dictd* servers making this knowledge easily available to everyone with an Web connection. Examine the available clients listed at www.dict.org/links.html to find one you prefer. Now that you understand the language, welcome to the group.


NETWORKS AND 64BIT COMPUTING

IBM's deskside servers

Big Blue is poised to grab a larger slice of the entry-level server market with a new two-tiered strategy to target potential Hewlett-Packard and Sun Microsystems buyers. The two new product lines will comprise a low-end Blade server range and an 'ultra low-end' rack system, both built around the Power PC 970 CPU which recently made its desktop debut in Apple Computer's new G5 machines.

The ultra low-end servers will run either AIX or Linux and prices start at \$3,500 for a four-way rack. Such aggressive pricing, IBM sources say, will target buyers investigating Intel-based commodity hardware.

IBM has also expanded its development effort to improve Linux on PPC hardware. Announcing the expansion of the Linux Technology Centre team from 250 to 300 programmers, Centre Manager Dan Frye said the company had not been



Linux on PPC is becoming more of a priority for IBM.

doing enough to perfect Linux on PPC.

"Linux runs pretty well today on Power," he said. "We want to take it from pretty good to world-class." As well as acquiring coders from other

areas of the company, Frye says new talent will also be sought to improve stability and hardware support. Efforts initially will be concentrated on high-end features including advances in

SMP facilities.

The move now puts all four major CPU manufacturers – IBM, Intel, AMD and Sun – firmly on the path to 64bit computing.

Linux Web Watch/



Online qualifications soon at MIT?



eserver.org has facts and fiction.



nationstates.net – Civ-ilised fighting.



Entertaining – craphound.com

Summer reads, make me feel fine...

From sneakers to studying, by way of Ulysses and usurping nation states!

Summer is traditionally a time when people tell you what to read on the beach. Well, there is no beach for LXF, so if you too are tied to your desk, try some of these small diversions.

Jennifer Government by Max Barry is a very fine satire on a future where branding is everything and the hero, Nike Hack, discovers a revolutionary way to sell sneakers. The book isn't online (you can read the first chapter at

www.maxberry.com), but a simple web-based nation-building game accompanies the text, and is at www.nationstates.net. Just don't mess with the United Socialist States of Salubrious!

Making good use of the Creative Commons license is Cory Doctorow, whose latest book, *Down and Out in the Magic Kingdom* is available in both deadtree and online editions. The

online publication can be read on the Net for downloaded for nothing. Get it at www.crashpound.com/down/. This is a sci-fi-ish satire on life, reality and the meaning of personal(ity) capital.

There is a ton of decent writing, mostly on academic subjects at <http://eserver.org>, but the project also hosts a wide selection of fiction, including short stories by the likes of GK. Chesterton and novels by Henry

James and James Joyce (including *Ulysses*, headache fans!).

And finally, if you're preparing for the beginning of the university year, you could do worse than visit MIT Open Courseware site at <http://ocw.mit.edu/index.html>, which features an ever-expanding range of research material covering everything from Aeronautics to Urban Planning. MIT claims it may soon offer online qualifications too.

NEWSBYTES

■ Netscape has release version 7.1 of its browser suite, complete with popup blocking and a host of other improvements. This is likely to be the last ever Netscape software release as AOL/Time Warner has laid off 50 developers working on the project. The Mozilla project, meanwhile, is itself going through a reorganisation phase, the result of which is a new non-profit entity dedicated to maintaining and developing Mozilla technology.



■ A group of Xbox/Linux hackers has released an exploit which cracks all of the consoles security systems, meaning Linux can be installed with no hardware modification. The group released details of the exploit after attempts to negotiate a 'signed' (ie official) Linux boot loader with Microsoft failed. MS hinted that it may pursue the group through the courts for 'facilitating piracy'.

■ Following the lead from Hewlett Packard (see right) Acer has announced it will begin shipping pre-loaded Linux PCs in Australia.

■ Mandrake developers are working to get the next version of their distro out by September. Version 9.2 may also herald Mandrake's return to non-protected financial status.

■ Wyse Technology has expanded its thin client offering utilising Linux with two new machines. Specs are at www.wyse.com/products/.

LINUX DTP Desktop debutante

Scribus, the native Linux desktop publishing application, has achieved its first milestone release. The newest version includes a number of useful additions such as support for MDI (multiple document interface), gradient fills, international hyphenation support, .pdf export (including support for transparency), CUPS support and a plugin for importing SVG format images. More importantly, **Scribus** has integrated support for both the CMYK colourspace and postscript output (essential for

Hewlett Packard's new Linux desktop machines are available online or from authorised resellers.

HARDWARE

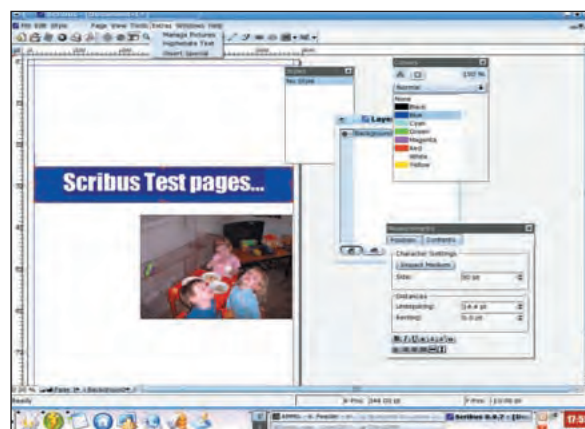
HP launches new Linux enterprise line

One of the world's biggest system builders has announced a new line of PCs aimed at small/medium enterprises preloaded with a choice of either Windows or Mandrake 9.1. The HP Compaq Business Desktop d220 Microtower features Pentium or Celeron processors with prices starting at \$349.00.

HP's Keith LeFebvre said the product was aimed at small businesses looking to maximise their IT investment without sacrificing quality or stability. "The combination of low purchase price and consistent, reliable performance in the HP Compaq d220 provides value-seeking businesses a solid platform choice."

The basic model (2GHz Celeron, 128MB/20GB, ethernet, Intel Extreme graphics) ships with Mandrake and, unlike previous Linux desktop offerings from major manufacturers (no matter how short lived), will pass on the cost savings to the user.

Technology commentators have suggested HP is emboldened by the ruling from Microsoft's adventures in the antitrust courts which prevents it from threatening or punishing vendors who have the temerity to offer non-Windows systems. It may also be the tipping point which brings Linux into the mainstream... Maybe Dixons and other UK retailers will stand up and be counted? www.hp.com



Scribus is a good, basic DTP package.

professional print work), and can use both TrueType and Type1 fonts.

Initial reviews suggest that Quark and Adobe don't yet have to wind up their

DTP divisions, but users of low-end applications from the likes of Microsoft and Serif may find **Scribus** to be a suitable replacement.

Jono Bacon

The founder of UK Linux, **KDE** developer and all-round nice guy, Jono Bacon is studying at Wolverhampton University.



COMMENT

Open or free?

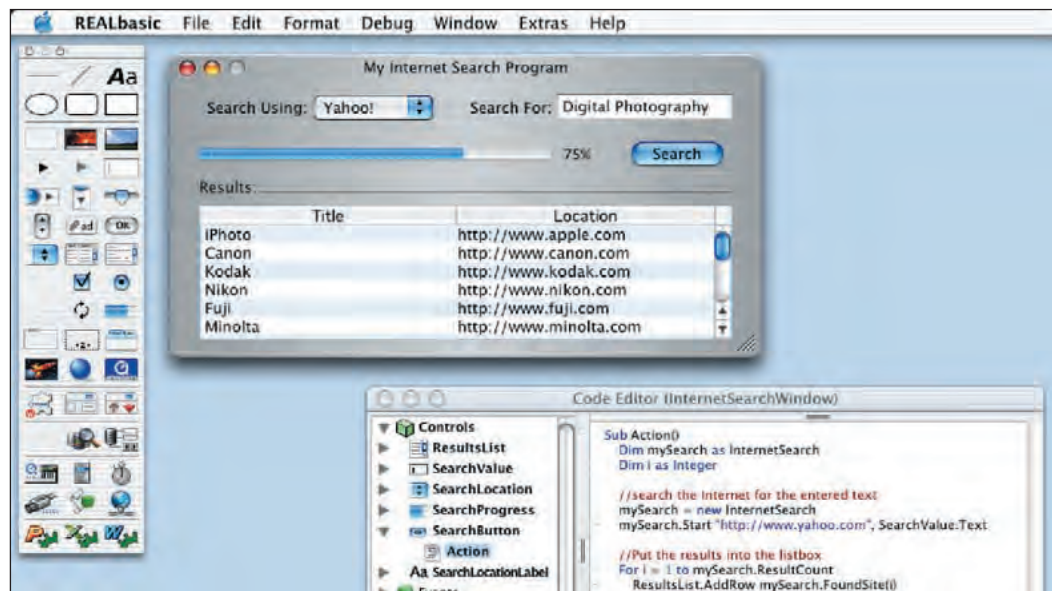
“ I had the pleasure of talking with Georg Greve from the FSF Europe recently. Georg was keen to stress the differences between Free Software and Open Source. This made me wonder if people really knew the differences.

Following up my discussion, I contacted Richard Stallman to discuss it further. Though I know that the aim of the FSF and GNU is to promote Free Software and the community that supports it, I wonder if the distinction that Open Source is a different concept is clearly made and as such, who has the responsibility of upholding the clear definition.

Many Linux folk think that the fact the source code is available is Open Source, and the fact that the software is free (in terms of cost) means that it is Free Software. The FSF and GNU work hard to show that free software is not just free in terms of money, but I feel the community aspect of the software is not necessarily drilled into the definition, and many Linux users simply assume that the community aspect of free software is a natural byproduct of distributed development. I think this difference between the two terms is an important one, and any confusion can undermine the fine work of the FSF and the OSI. I have a lot of respect for these organisations in upholding the ideals and goals that surround free software and Open Source. It is important that users are correct in their usage of the terms. I myself have been guilty in misusing the terms before, and although I am not a hardcore FSF/GNU/OSI member, I'm actively involved with the community; if there is confusion for myself who has been doing this for years, how can new users be expected to know the difference? ”

COMPILER/IDE

REALBasic ready for Linux?



Mac and Windows developers can now compile REALBasic applications for Linux.

Real Software has announced support in the upcoming release of REALBasic 5.5 for Linux as a target environment, and also hinted that a future version would be a fully fledged host for the IDE. Real's Head of Marketing told The Register that compiling for Linux from within the Macintosh or Windows environment

would meet the needs of most of their customers for the immediate future. CEO Geoff Perlman said adding support for Linux was an evolutionary step. "Our vision for REALbasic has always been to help people create software more easily, regardless of their computing platform," he said.

Long-term Real user Dan Farrand

of Green River Computing said the addition of Linux support makes it the only high-productivity rapid application development tool for all three platforms. "From a single code base, REALbasic gives us the power and flexibility to provide our software on the platform our customers choose, and Linux is what we're hearing."

SCO News

■ SCO is preparing a new Linux licensing program, which, it says, will allow buyers to continue using Linux without fear of litigation. Spokesman Blake Stowell said SCO was keen to offer a way for people 'to be able to run Linux legally.' Some commentators have suggested that this move could be seen as a sign of weakness or a last ditch attempt to raise revenue from the current litigation against IBM before the case comes to court.

■ Transcripts from former SCO CEO Ransom Love's keynote speech to the LinuxWorld conference in 2000 suggest that Caldera (fresh from its acquisition of SCO) would be putting code into Linux to ensure it scaled up for Intel's IA-64 processor. Responding to a question about Project Monterey and its potential for a clash with the IA-64 Linux development effort, Love said: "Clearly we are going to add components back to the Linux kernel on both IA-32 and IA-64 platforms. We'll work with Linus and everyone in order to make that available." More information is available at <http://twiki.iwethy.org/twiki/bin/view/Main/TrillianProject>

■ Darl McBride, CEO of SCO, dashed off to Japan recently after a consortium of the world's biggest electronics manufacturers joined forces to develop Linux for home entertainment devices (see ELN News, below left). McBride, a fluent Japanese speaker, was said to be anxious to put SCO's case before the consortium committed itself to Open Source development.

Embedded Linux News

● After a short campaign, **Linksys** has released the code which lies at the heart of its WRT54G Wireless Router. The release of the various elements which power the device was made to keep within the spirit and letter of the GPL. The code parts are available from www.linksys.com/support/gpl.asp.

● **Real Networks** continues its program of opening technology to developers with the release of the source code to its Synchronized Multimedia Integration Language (SMIL).

SMIL can be used to synchronise web content or create individualised presentation that is dependent on criteria such as language or available bandwidth. If you want to see how the project is maintained, along with the rest of Real's other OSS efforts, visit www.herlixcommunity.org.

● Eight of the world's top electronic manufacturers have banded together to develop the Linux OS in the Consumer Electronics (CE) sector. The **CE Linux Forum** (CELF) counts Matsushita (Panasonic), NEC, Philips, Samsung, Sharp, Sony, and Toshiba as its founding members and has already set up its first working group looking at specific requirements for efficient CE use. The group has said that code developed under the project will comply fully with the GNU GPL and, as such, the process or creating specifications and implementations would also be a public process. CELF builds on the existing partnership Sony and Matsushita established last year to create extensions to the OS to make it more suitable for devices such as digital camcorders and set top boxes. Rumours suggest that IBM will be seeking to join this select group in the near future.


MULTIPLE OS CAPABILITY

VMware upgrades ESX Server

Virtualisation specialist **VMware** has updated its flagship ESX application which brings multiple operating systems to high-end systems and blade servers from the likes of IBM and HP. The new release also features an upgrade option called Virtual SMP that can create a virtual dual-processor machine on a single processor server. A four-processor version is slated for release next year.

In addition to virtual SMP technology, VMware will also debut a new management tool to move server

Virtual SMP will bring dual-cpu applications to single processor servers.



applications from physical to virtual servers. The P2V application complements the company's recently announced *Control Centre* application.

The ESX package costs \$3,750 for a dual-processor license, while the Virtual SMP add-on costs \$1,250. Price doubles for a four CPU version and doubles again for an eight-way package.

LINARE

Another new distro



The Linux Desktop market is set to ever more crowded with the launch of a new hardware and software package from Linare, a global company intent on bringing the joys of Open Source to the masses. Following the Windows

paradigm, Linare is offering its operating system, built around KDE, *OpenOffice.org* and *Mozilla*, as a single product (priced at \$19.99) or pre-installed on a basic monitorless PC for a very competitive \$199.99. The

hardware specs compare well with other low-end packages, featuring a 1GHz Via processor, 128MB RAM and 20GB hard disk.

Both products are expected to be available from August 2003.

SUSE HARDWARE

Wal-Mart expands Linux offerings

American retailer Wal-Mart has added SuSE 8.2 as an option to its range of low-cost PCs. Buyers can now choose between a traditional Windows-based PC or opt for a machine preloaded with Lycoris, Mandrake or SuSE. Prices for the SuSE package, still only available online for US buyers, start at \$298. The most basic product includes a Duron processor, 128MB RAM, 20GB hard disk, CD ROM, integrated ethernet and modem, while more expensive Athlon-based options include more memory, bigger disks and a CD/RW drive. Holger Dyroff, SuSE's US General Manager, said the expansion of choices reflected "the growing demand for affordable and innovative Linux systems and the widespread adoption of open source."

LESS PACKAGING!

Red Hat unboxed

Red Hat has decided to alter the way it markets its distro to general users. The effect of this move is that you will no longer see boxed products on computer store shelves. The company has said the next release, codenamed 'Cambridge', will not be sold through traditional retail channels, allowing the company to focus on its revenue making support business.

In tandem with the withdrawal of the boxed product, Red Hat says it will in future make development lists public and hand back package maintenance to application developers, instead of undertaking the work in-house.



Red Hat boxes will soon be a thing of the past.

A spokesman said both moves were designed to reduce distribution lead times which often mean applications are 'out-of-date' before the CDs have been duplicated and boxes printed.

David Cartwright

David Cartwright is an IT consultant who specialises in providing Linux systems and solutions.



COMMENT

TLC from DLC

In the last few days, the Desktop Linux Consortium www.desktoplinuxconsortium.org/ has winked into existence. I do hope it achieves its stated goal: "To promote the wide scale understanding and adoption of the Linux OS for use on the desktop." But they face an uphill struggle.

There are a number of problems with supporting Linux on the desktop in the real world that need to be addressed if it's going to become popular. Many people I've asked see lack of app support as the main issue, but I disagree – there are actually plenty of apps out there covering most of the 'common' desktop requirements.

In fact, I see the problem as lack of support from the manufacturers of hardware and peripherals. The main problem I've always had with my various PCs is that they don't support Linux for a few weeks after I've acquired them. They ship with Windows 98, 2000, XP or whatever, and with a few minor exceptions (such as my Samsung printer/fax that XP doesn't understand) all the hardware works with the operating system. Compare this with Linux – when I've bought a new PC I've had to wait for some update or other before Linux could make (say) the sound card work, or X-Windows understand the graphics adaptor.

OK, some vendors are now formally supporting Linux, but they're in a minority and while this continues, Linux doesn't stand a chance of desktop support (who would use an OS that only works on last month's technology?). My hope is that the DLC will quickly begin to address this issue with the vendors – and once the big names have bought in, I do actually think Linux stands a very good chance on desktop computers.

Mailserver

Share your opinions, right wrongs and demand justice by writing to *Linux Format*. Drop us a line at: **Linux Format**, Future Publishing, 30 Monmouth Street, Bath BA1 2BW or email: lxformat@futurenet.co.uk

★ Letter of the month

This month's winner receives a copy of **Robot Building For Beginners** (ISBN 1-8931-1544-5)

Legally Linux

I work as an operator (Network printing etc) for a large law firm in the City of London. Recently we replaced our aged PrismExpress Scanning system with an A0 scanner, Dell rip server and a Xerox A0 printer. I was so surprised when the machine was set up – it's got Linux at the heart of it running Mandrake 8.2! (I have been a Mandrake user since 7.2, currently running MDK 9.1 on my Compaq laptop). Even better, the machine itself is actually a AMD XP1800 DDR333 machine, so it breaks two myths at the same time about AMD and Linux not being ready for the Enterprise market!

This little penguin can handle just about anything we throw at it. A0-size full-colour plans are scanned in about 10 seconds, 1 minute to rip it and 3 minutes for the print. The printing speed has nothing to do with Linux (the Xerox takes the credit) but without the penguin to handle everything, it would not be possible at all. Our old system took about 15 minutes to do just one! On top of that, the app is developed by... Caldera!



Mandrake 8.2 is the power behind this commercial print setup.

I have never seen anything like this before. It not only comes with a full-featured network listener and Spool controller, it also comes with two full-featured editing utils (I suspect it's *The GIMP* at the core). To my surprise, lots of people learn to use the Linux machine (well, the app) in about an hour even though a lot of them have never had any computer training.

Conclusion – Linux is very much ready for Enterprise!
Joel, London

Did you ever doubt it? Linux is pretty pervasive in all sorts of 'black box' applications, though I admit it is a surprise to find it running on Mandrake. For your keen observational skills, and to help you get started building other Linux devices of your own, please accept the star prize this issue – *Robot Building for Beginners*. If any other readers have come across Linux running in places they didn't expect, or on hardware that you probably wouldn't immediately associate with our favourite OS, please let us know!

Backtracking

A few months back I switched from Mandrake 8.0 to 9.0. Big mistake! Annoying problems with two different PS2 wheel mice (Logitech and Microsoft), and after several re-installs to find a fix, I lost *drakconf*!

I found your June issue at the local mall (a scant 5500 miles WSW of you) and upgraded to 9.1. Everything is fine with the mouse

and *drakconf* is restored, but now, despite correct hardware identification, there are problems with the modem, printer and graphics configurations. I went back to 9.0, which I can live with. Next stop, SuSE 8.2!

Jim Wilson, Diamond Bar, CA, USA.

In spite of great advances in Linux installers over the years, there are occasional bits of hardware that seem to cause problems. On a different

computer, you may have found the install to go smoothly. Usually these problems are ironed out in time, but of course, by then there is new hardware to cause brand new problems!

Helpdex howler

Hate to p*s on your fire ... But I noticed a small mistake on your second cartoon in *LXF42*...

The last frame should surely read ...

"... A Supernova's the celestial equivalent of "rm -fr /" with root permission"

not

"rm -fr \\"

Or have we been using Windows for too long ?

Iain, via email

I'm sure it was a mere slip of the pencil on Shane's behalf. Thanks for the correction though.

Sorted by LXF!

Thank you for publishing and replying to my letter, (headed "Time tinkering", *LXF42*). I thought I would let you know how I was getting on. I solved the problem with the Nvidia drivers and the Soundblaster card. You were of course perfectly correct, it was a hardware problem, but I couldn't have fixed it without the help of the Linux community.

As I had mentioned, I had been lurking on Nvidia's Linux forum, so I kept going back and refining my searches for KT400 postings. There were quite a few fixes which I tried but without success. Finally I came across a posting that suggested disabling APIC (Advanced programmable Interrupt Controller) in the BIOS. This fixed the problem with both the Soundblaster and Nvidia drivers, Linux now works really well and I am really looking forward to the day I can finally dump Microsoft Windows forever!

It took me about two months on-and-off to fix it, but with my searching I did find out how to set the AGP drivers properly, which I wouldn't have done if everything had worked correctly in the first place. It has also taught me that the best place to sort out a problem is the Linux community itself. Thanks for a great magazine.

R. Sutcliffe, via email

READER TIPS

LINUX DIALOUT

I thought I would tell you how I solved the MDK missing dialout problem. The 2 CD set of Mandrake 9.1 is missing the KPPP executable, plus also the rules and icon files. I solved this by copying the directories from an old installation of Mandrake 9, out of the Home directory, and the executable from /usr/sbin, if I recall correctly. Then copying them to the new installation of 9.1, add the menu entry and away you go! The actual distro is pretty good in my opinion, lacking only ready-installed Flash and shockwave plugins for Mozilla.

Bryan, *via email*

The files are available on any of the Mandrake mirror sites. Sorry to anyone who had difficulties – we're not sure why Mandrake left them out of the distribution set to start with.

MDK PREJUDICE?

In the letter 'KDE Screensavers' in LXF issue 43, Dave Spagnol said "As usual, Mandrake overwrites things behind your back!" The practice he refers to is caused by setting higher security levels and accepting default settings for them, which causes msec to return a number of things to what is considered more secure. All these things are configurable, but it does take effort to learn about it. May I suggest that MDK novices:

1 Join the newbie mailing list – this can be done from Mandrake's website. It is very active, so be prepared for a lot of mail, but you will learn a lot from it.

2 Read relevant pages from Mandrake's website – searchable, of course!

3 Visit the Community TWiki site, particularly the page <http://mandrake.vmlinux.ca/bin/view/Main/NewbieFriendly>

The MandrakeReferences page on the same TWiki has lots of links for more experienced users.

Anne Wilson, *via email*

BORN TO BLEND...

I am a fledgling 3D artist with a very shallow wallet. Blender had me on the price range but nearly lost me on the learning curve and the lack of decent info didn't do much to resolve that. However, I your tutorials are an invaluable resource where so few are present. I would only have two resources to add: most Lightwave tutorials can be adapted for use in Blender and irc.debian.org has several good channels for Blender discussion including #blender and #gameblender for discussion of the Blender Game Engine.

Rob, *Independence, MO, USA*

BELKIN SWITCH

Re: 'SuSE slipping', Answers section, page 95 LXF43 – I had the same issue using my Belkin switchbox. Until I found the fix, I had a workaround – switch to a console and back: (Ctrl-Alt-F1, Ctrl-Alt-F7 – IIRC)

However, I finally got fed up and searched the web, eventually finding a link into the Belkin support site, where they suggest changing to 'IMPS/2' – In practice, I fired up YaST and picked the first

option – PS/2 mouse aux port. Ben Prescott, *via email*

WINDOWS? NAH!

Having read many tales of woe in the letters pages regarding Linux compatibility I thought it might be worth injecting a little optimism. I use Slackware 8.1, X drives my SiS graphics card at 1152 x 864, a resolution that the manufacturer's docs don't mention but is as much as a 15-inch monitor (and my eyes) can cope with. It also handles a generic serial mouse and a Logitech trackball, simultaneously. ALSA deals with both my Soundblaster Live and SB16 giving Rosegarden five synths on the SB Live and one on the SB16 plus two MIDI ins and two MIDI outs, Smurf takes care of the soundfonts. Gphoto2 gives me access to a cheap and cheerful USB digital camera.

On the SCSI bus I have a SyQuest removable disc drive, a DAT drive and a CD-R. Vuescan handles an Epson flatbed scanner and a Canon film scanner. The GIMP does the photo post-processing and via gimp-print drives the Epson 2100 A3+ inkjet with excellent results. I used to use a serial modem until I needed to port to talk to a old Psion I picked up. It's now on the firewall machine (Linux, of course). OpenOffice.org has got me through two Open University courses despite the OU demanding I use MS Office (very short sighted I think). Oh and Konqueror allows me to use Internet banking despite pop-up windows instructing me to use a compatible browser. All this may be possible under Windows; I don't know, I haven't had to find out. Andy Horseman, *via email*

Well done for your perseverance. I haven't heard of BIOS APIC settings causing problems with either sound or graphics cards before! This is something we'll have to bear in mind for future problems we might receive.

As well as our Answers pages, you'll find the Linux Format forums a great place for help and advice. Check them out at www.linuxformat.co.uk (click on the link to the forum in the left-hand menu). Even if the throng of Linux fans there can't answer your question directly, they can usually give

some helpful advice as to possible causes and where you might be able to get an answer.

Linux out of the box

Having tried to see what Linux was all about several years ago, with no success I might add, I put the distros away. Being a beginner I didn't know how to configure them – I was used to being babied by Windows. The distros then were Mandrake 7.2, Red Hat 5.2, and Caldera 2.2. Last Christmas I

spotted your magazine on a local news stand. It had Mandrake 9.0 with it. This stirred the longing for trying something different than the blue screens of death that I was living with. I now have Mandrake 9.1 on a desktop computer and Red Hat 9.0 on my IBM Thinkpad. And I am able to print and surf the net. That is what I mean by getting Linux out of the box. Before, I couldn't get it configured for Internet or printer. It was, as far I as I was concerned, stuck >>

UNIHOST

Specialists in hosting dynamic web sites & e-commerce stores with secure SSL.

- Linux Hosting Specialists
- CGI / SSI
- PHP 4
- Perl / DBI
- MySQL
- PostgreSQL
- JSP / Java Servlets
- FTP / SSH Access
- Full Logs & Web Stats
- Huge Network Backbone
- 100 MB Storage
- Full Shell Access
- Monitored 24 x 7 x 365
- UPS
- POP3 E-mail Boxes
- Unlimited Aliases
- SSL Option Available



UNLIMITED BANDWIDTH

£19.99 per month

SPECIAL OFFER:

20% off 1st year's hosting if booked before Sept 30, 2003. (£16.99/month) Quote ref: LF903M

sales@unihost.net

0870 744 0721
UNIHOST.NET

Mailserver

« in the box or computer case. I had no way of using it until I could set it free.

Bobby Lyle, USA

All the major Linux distributions are more accessible these days. The problem used to be the installer, which often necessitated knowing the ins and outs of your hardware very well. That has been pretty much solved in the last year or so, although brand new hardware and peripherals will probably always be a problem.

With the install solved, the post install has become more a centre of attention – networking, printing, filesharing and so on are easier than they were, but by no means transparent enough for many users.

MySQL bias? Us?!

I very much enjoy reading your magazine. It is certainly in a class all its own. However, I have been moved to write due to my annoyance at the way your magazine presents the DBMS: MySQL. This culminated last issue when a two-page review was given to the latest version.

Specifically it was rated 9/10 for features. This is, as even MySQL will admit, is just not true. MySQL is very low on features. MySQL has a bias towards performance and not features. MySQL doesn't have

Views, Triggers, Functions, Sub-queries or Check Constraints. All of which are part of ANSI SQL and used in most other DBMSes. Eg Microsoft, PostgreSQL, Oracle and most others anybody would care to mention. It is wholly accurate to mark it 10/10 for performance, but please be accurate, mark MySQL 3/10 for features.

This is a very important point. MySQL says in its documentation that features like Views, Procedures and Triggers are not required much. This is a vast exaggeration of the truth. As a DBA, not having these features is a chronic constraint. For instance, Views are extensively used in CGI applications to protect the database from hackers, bugs and other problems, lower the development time, and give a better solution – as suggested should be done in any Database theory book you care to read. I could give as valid examples for the others I mention. These are not rarely used obscure toys, but the very heart of a modern DBMS.

Your review also failed to mention the licence issue with MySQL, which is GPL. If I care to use PostgreSQL, (under BSD licence) I can program with it at work for my commercial applications, and at home I can

program the DBMS itself without issue. This I cannot do with MySQL due to their highly aggressive licensing policy. Here I would pay to use it at work, and as their coders are largely in-house and on salary, I would be unlikely to have a patch accepted for the product either.

People compare MySQL to Linux, and quote, (as did Tim O'Reilly at City University recently) that MySQL will do for DBMS what Linux did for the OS. This misses out an important fact. Linux

the system, shown in the rather bizarre and hard to fathom error messages InnoDB gives out, and it failure to adhere to ANSI SQL. (Eg using '1' and '0' for Boolean, instead of the written standard stating 'TRUE' and 'FALSE' should be used. MySQL 4 doesn't even recognise 'TRUE' and 'FALSE'.)

All in all, I wonder whether your magazine is showing undue bias towards this product. For instance, I await eagerly for as glowing two-page review about three other far

“While it's true to say the BSD licence is the 'most free' in terms of the rights it gives you to the code, I believe that the GPL is better at encouraging freedom.”

Torvalds was against the commercial licensing of Linux from the start. I wonder if Linux would have successes if Linux was as aggressive about the use of commercial licences for Linux as MySQL are?

There are other issues connected with MySQL also not mentioned, like the immaturity of

better Open Source solutions, PostgreSQL, SAP-DB and Firebird; all of which are more open, more mature and have more features.

Ben Clewett, via email

Thank you for your letter. I'm probably going to surprise you by agreeing, at least in part. Compared to many other SQL implementations, MySQL is light on features, mainly for the

Helpdex

shane_collinge@yahoo.com





Linux distros have come a long way since LXF started, and installs get easier with each new iteration.

sake of speed. I think 3 would be a bit harsh, but perhaps 7 might have been fairer. All reviews are to a large extent opinion though..

I do have to take issue with the licensing argument though. We had a bit of a debate about this in the office recently, and while in the strictest sense it is true to say that the BSD licence is the 'most free' in terms of the rights it gives you to the code, I believe that the GPL is better at encouraging freedom. The GPL was designed not only to give the 'user' the freedom to do what they wanted with the code, but also to protect that freedom for other users too. If you don't agree with the terms of the GPL and wish to develop commercial applications based on the code, which you then distribute in a proprietary fashion, I don't think it's unfair to have to choose a different, paid for, licence. Some projects may choose to use the less restrictive licences such as BSD, but I really don't think that the developers of MySQL can be accused of being unfair, or even of being particularly aggressive.

PostgreSQL is to be the subject of a forthcoming feature in the magazine, and also make sure you don't miss LXF's next issue where we will be examining FreeBSD.

Ready for desktop?

Having used Linux, but under the admission that I'm hardly a guru, for several years I have seen the views both sides of the 'Is Linux ready for the desktop?' debate. The main argument against Linux appears to be that it is hard to set up and some things do not work if

you do not know how to fix them. For me, I'd rather know that there is a way to fix the problem than sit in front of my computer cursing at control panel for not giving me enough control. There are many tools available (for example in the Mandrake distributions) that will mimic the Windows control panel well, but it is comforting to know that below this there is often even more control and tweaking that can be done.

People often seem to forget that Windows is less than perfect. This isn't a rant on the subject of how useless Windows is because in reality it's quite good, however people forget that sometimes Windows doesn't always work as it should. I have sat for countless hours pondering why my computer won't work with the same setting as the last time I installed the hardware/software, exasperatedly wondering why my computer is doing what it is and not giving me a choice in the matter.

I think Linux has come on in leaps and bounds since I started using it. There is very little that I am lacking on my Linux box now. When I hear people talking about how hard Linux is to manage I start wishing I could see them next time their Windows box goes wrong.

PS. relating to LXF42, 'The Subtle Flavours of WINE', all software is created equal, it's just that some are more equal than others.

S Jeapes, Imperial College London

I think there is a certain amount of rosy-spectacledness going on sometimes when people compare the relative ease of use of OSes. [LXF](#)

SUBMISSION ADVICE

WHAT WE WANT:

- Letters about the magazine or Linux in general
- Constructive criticism
- Your opinions
- Concise points about relevant subjects

WHAT WE DON'T WANT:

- Technical question – direct those to our Q&A pages!
- Random abuse
- Nonsense rants
- 200 pages of meandering diatribe

WRITE TO US AT:

Linux Format, Future Publishing, 30 Monmouth Street, Bath BA1 2BW or email: lxfl.letters@futurenet.co.uk



»There is no reason for any individual to have a computer in their home.«

Ken Olson, former President, Digital Equipment Corp., 1977

»Linux is only for computer science students.«

Another myth disproved !



SuSE LINUX is for everyone – anytime – anywhere!

SuSE Linux requires no prior training or special hardware. Simply install and start using a multitude of applications that allow you to work and play as you wish. An ideal replacement for your old operating system, whether on your desktop PC or your laptop, it is just as easy – but faster, more secure, and much more economical.

SuSE LINUX 8.2 PERSONAL

- The Concise Linux for beginners and home users
- Easy working with OpenOffice suite (compatible with MS Office)
- Games and handwriting recognition programmes
- Multimedia applications (e.g. CD recording, MP3 and MoviePlayer)
- Secure e-mail and Internet access

SuSE LINUX 8.2 PROFESSIONAL

- The Complete Linux for advanced and SOHO users
- Automatic network identification for your mobility
- Multimedia and video editing to professional standard
- Optimised installation, operation, and help functions
- Numerous development and programming tools for the experienced user

@ info@SuSE.co.uk

+44(0)208 846 3918

+44(0)208 748 0538

SuSE LINUX Ltd.
Appleton House
139 King Street
Hammersmith
London W6 9JG · UK



Reviews

All the latest software and hardware reviewed and rated by our experts

LXF VERDICT EXPLAINED

Each review is accompanied by a Linux Format Verdict to help you to assess the product at a glance (it's no substitute for actually reading the review, though). We award scores out of ten in the following categories:

Features: Does it provide the functions you need? Is it innovative?

Performance: How well does it do its job? Is it fast and reliable?

Ease-of-use: Is the interface well designed? Is the documentation well written, helpful?

Value for money/Documentation: Whichever is most appropriate!

For those who like numbers, the *Linux Format* Rating is a score out of 10 summing up the overall excellence of a product. It will usually, but need not be, an average of the above categories. We award scores as follows:



10 The close-to-perfect product.



8-9 Good, but has a few niggles.



6-7 Does the job, but needs work.



5-4 Average.



1-3 An utter disaster. Back to the drawing board.

THE TOP STUFF AWARD

If we really, really like something – we really think that a particular piece of software, hardware or any other sort of ware is the best stuff around – then we'll give it our *Top Stuff* Award. Only the very best will be chosen. It's not guaranteed to all products that score highly.



WHAT'S NEW...

Neoware Eon >>> Preferred 4000

Thin client computing is making a big comeback as Linux makes it easy to read data from a remote server **p18**

Seapine SCM

Command-line CVS frustrating you? Try this out! **p20**

Borland JBuilder 9 Enterprise

Java development from the folks famed for *Kylix* **p22**

SuSE Enterprise Desktop 1

A great option for both users and admins migrating from the Windows environment **p24**

Ximian Desktop 2

Integrated desktop that quite literally rolls out the *Red Carpet* to make your life easier **p26**

Navaho TeamCAT

A server that looks smart enough to impress your mother and will amaze your boss with its performance **p29**

Books

Java Pitfalls, *Python In A Nutshell*, *Programming C#*, *FreeBSD Unleashed* (2nd Edition) **p30**



LINUX FORMAT BENCHMARKS EXPLAINED

To provide objective performance comparison between machines running Linux, we run a set of in-house benchmarks. These are: *bonnie* and *hdparm* to test hard drive performance ('hd' in the benchmarks), MySQL *Super-Smack* to test how well a machine handles database serving ('mysql'), *ApacheBench* to test how fast a machine can serve web pages ('apache'), a *gcc* compilation of Linux kernel 2.4.19 ("compile"), and *oggenc* to convert a test .wav file to a .ogg file. These numbers are then averaged to produce an overall score, which may be adjusted

slightly now and then, if a machine has a particular high or low point that should be taken in to consideration. Combined, these tests really push hard drives, network cards, and CPUs to their limits, and so give quite a representative figure – a multiple of the performance our yardstick machine.

The LXF yardstick box attempts to represent an 'average' reader's setup: Debian 3.0 on an 866MHz PIII with 256MB of RAM. So, a machine which scores 1.5 on our *Apache* test served 50% more web pages than our yardstick, whereas a box that scores 3.0 for overall ran, on average, three times faster than our yardstick box.

BENCHMARKS

hd:	0.83
apache:	1.22
mysql:	1.11
compile:	0.96
oggenc:	1.71
Overall:	1.17

The blue bar in the example above represents the performance figure for the hardware, and the red bar is the benchmark figure. When a piece of kit performs lower than the benchmark, as in 'hd' and 'compile' above, the blue value will appear less than the red value. **LXF**

THIN CLIENT LINUX APPLIANCE

Neoware Eon Preferred 4000

20 years ago, thin clients were the 'in thing'. Now they're back, and in style – but are they better than ever? **Paul Hudson** tries one out, and finds himself listening to the sound of silence...

BUYER INFO

There's lots of info on thin client Linux on the Web, try www.lts.org and www.citrix.com for starters.

- **SUPPLIER** Neoware
- **WEB** www.neoware.com/products/eon/preferred.html
- **PRICE** \$569

The debate on thin clients has raged for longer than most people can remember. The question is simple: should workstations hold their operating system and applications locally, or should the data be read from a remote server? Both methods have been favoured in the past – thin clients were particularly popular in the days when Unix was in its prime, and hundreds of people would connect through dumb terminals to a mainframe. Then, in the early 1990s, Windows rose to desktop dominance, changing the paradigm so that PCs worked standalone, utilising to the full the power of the new CPUs Intel was producing at the time.

Mostly thanks to the rate that PC hardware improved, it's been some time since the question of thin clients was reconsidered. However now that almost every company relies on administering its own computers to succeed, more and more admins are looking to solve the problem of distributed administration. With every PC being unique and different, it's incredibly hard to monitor software installations and patches, meaning that administrators have to be scattered across a corporate network to make sure they can be on-hand whenever there's a problem.

Neoware, a company we looked at in issue 41's *Linux Pro*, has a powerful range of thin clients appliances,

several of which are based on Linux. Neoware sent us an Eon 4000 model, which is an unassuming, grey-ish box with no CDROM drive, and two USB ports on the back. The lack of a CDROM might seem problematic at first, but that's nullified by the presence of USB – you can just plug in devices as you need. This new wave of thin clients are designed to do everything PCs do, without the administration worries.

The Neoware solution

We plugged ours in and turned it on to find... nothing. No fans, no whirring, no beeping, no sound at all. And yet all the lights were on, and I could clearly see it was booting up as it output messages to the monitor – this thing is silent, and that's not silent as in "very quiet", it's literally absolutely noiseless. A whole room of these would make less noise than just one conventional PC, which is incredible.

Inside there's a 300MHz National Semiconductor Geode CPU, which is Pentium MMX-compatible and runs quite quickly. There's also a basic on-board video card, and an old Linux kernel and XFree86 version (2.2 and 3.3 respectively). Nevertheless, the machine boots up into X in under a minute, and you're into the basic Neoware connection management system, *ezConnect*. This tool, which has a friendly, graphical interface, allows you to configure your Neoware machines to connect to terminal servers, and also to set the machine up to always use a given connection – effectively making the box appear to be a thick client.

Neoware has got a lot of experience designing and producing thin clients, and it shows. The Eon is designed to work as an appliance in the same way your refrigerator is – you plug it in, turn



Pros and cons of thin clients

PROS

- Usually Cheaper to implement
- Reduced IT management overhead
- More effective use of resources
- Scalable
- Various remote admin options

CONS

- Not suitable for some high performance tasks
- Limited usability of external devices
- High server/network demand
- Can introduce a single point of failure – eg if the network/server goes down

it on, and forget about it because you know it will just work.

Configure the system

By default, *ezConnect* allows you to connect directly to a Citrix MetaFrame server, but it can also make simple connections to Linux machines. The simple connections are, sadly, very simple – it uses telnet (which we haven't used in years), with SSH not even being an option. One feature that is very noticeably missing is the ability to run a remote X-Windows session on the machine. While it's probably possible to hand-reconfigure the unit to connect to a remote system, that wouldn't make sense given that it's supposed to be easy to set up and use. If *ezConnect* was extended to include seamless support for connecting to other Unix machines, the machine would be a lot more portable. As it stands, if you don't use Citrix/Windows, you're basically out of luck.

However, the GUI is remarkably simple – you can perform all the various actions you'd need (setting up networking, display information, rebooting, getting access to a terminal, etc) using your mouse. It's so easy to use that we had the machine connecting perfectly in under three minutes, despite not having a manual to refer to.

Which brings us to one minor hitch: there are no manuals included in the box, just an extremely basic two-page connection guide and warranty. No "What's in the box", or a recovery CD. It's possible that this sorry state of affairs is just because we have a review machine – at least, that's what we're hoping. However, we only consider it a minor hitch on the grounds that if you're purchasing a hundred of these

things to work on your network, it's not likely you'll be wanting to pay the cost of having a hundred manuals. Furthermore, it's hardly like these machines are hard to administer – indeed, they're more like an administrator's dream come true. The lack of a recovery CD is inexcusable, though; particularly because the software that comes on the box isn't readily accessible elsewhere – it appears to be a very stripped-down version of Red Hat.

Connecting to a server is a doddle – after using the childish simple GUI that hand-holds you through connecting to a Citrix server, you're off and running straight away. Even though the internal hardware would have been considered old even four years ago, it performs very quickly across the board – simply because it isn't doing much work at all other than displaying graphics. The hard stuff is being done elsewhere.

Using the system

From the user's point of view, the system might just as well be a normal PC – it's just as fast and powerful, without the noise, size, or instability. That is, not only is the Eon configured and hosted remotely, but everything that is local is locked away from users – they have limited privileges even for the local home directory, and the filesystems are all marked read-only so that unexpected power outages have no effect on overall stability.

In environments where thin clients are likely popular – schools and universities, and, to a lesser extent, businesses – it's likely to be considered a downside that there's the need to plug in devices by USB. In fact, unless the devices are somehow nailed down and away from users' attention, there's the chance that having detachable peripherals might end up being more of a hassle for administrators than just using normal PCs.

Ups and downs

Switching back to the paradigm of thin clients comes with a bundle of pros and cons that need to be weighed up before any commitment is made. For example, despite Neoware's machine being ultra-quiet and inexpensive, with just 300MHz powering it, it's never going to be able to do any serious processing tasks by itself. While it's possible to shift a lot of work onto the server, it's still not possible to run CAD/CAM software with any degree of

If you really need removable media, there is the usual selection of connectivity at the back. That there's only two USB ports shouldn't matter, as this box is primarily a network solution.

success, or even to play modern games. Another big problem with having your network built around thin clients is that you need a fairly hefty network backbone – with virtually all important data files stored in a server somewhere, networks easily get clogged as the data is transferred to and from clients..

One minor point about these machines is that if the central server were ever to go down, each and every thin client depending on it would also become useless – they have no means to operate by themselves. However, these flaws are inherent to thin client computing, not the Neoware machine specifically, so we can't really hold it against the box – however, it's important to keep in mind.

Conclusion

On its own merits, the Eon has little against it – as thin clients go, this does all you need. While it would be nice to have a CDROM built in, it just adds another point of failure. Without a fan or indeed any many of the moving parts other PCs have, this thing really has just one or two points of failure – either the network card or the CPU will fail, both of which are very unlikely.

As such, this machine is everything it's made out to be: fast, very easy, quiet, and hardly noticeable. It's smaller than your average satellite transceiver box, which means you can hide them anywhere. It's not perfect – we'd like to see the GUI extended to allow you to connect to other X Windows machines for your session, for example, but that will hopefully come in the future. In the meantime, it's a neat little Linux appliance that's available at a great price. **LXF**

VERDICT

Features	9/10
Performance	8/10
Ease of use	9/10
Value for money	8/10

Tinier than tiny, but still able to perform in today's business or academic environment.

LINUX FORMAT RATING
 **8/10**



SOURCE CODE MANAGEMENT

Seapine Surround SCM

This source code management server and client touts itself as a system that can “Manage code, manage change – anywhere, anytime”. Strong claims, but does it live up to expectations, asks **Jono Bacon**.

BUYER INFO

The the command-line nature of CVS doesn't appeal, this could be your ideal application.

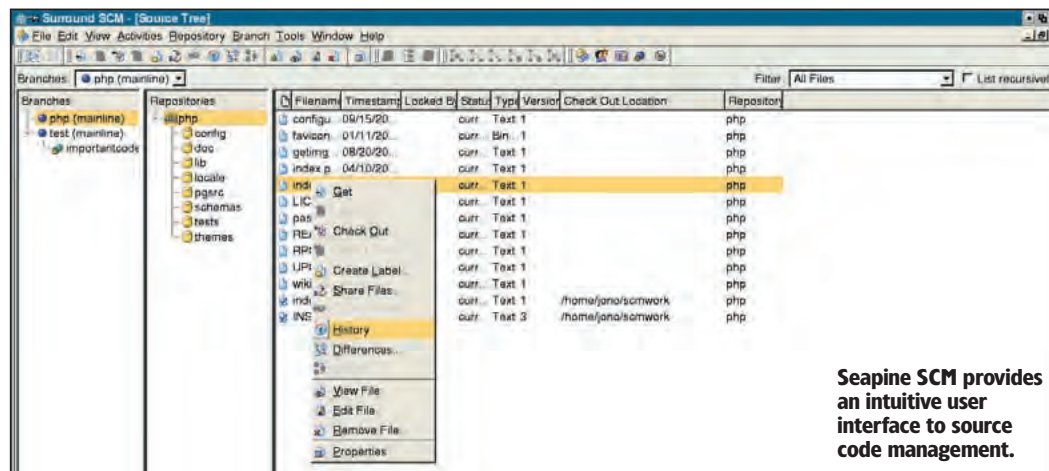
- **DEVELOPER** Seapine
- **PRICE** from \$595
- **WEB** www.seapine.com

Source code management is generally a somewhat uninteresting, yet important aspect of software development that can get overlooked by many in the Linux world. *Seapine Surround SCM (Software Change Management)* is a true example of how software should be installed. A simple Java-based GUI installer was used to install the system, asking where to install files and other data. After installation was complete, it was just a case of running the server and then running the client. The program ran without a hitch, first allowing me to log in with the default password, and then even suggested that I should think about changing it immediately!

In use

Seapine SCM has a rather impressive featureset including facilities for adding, removing and updating files, viewing history info, dealing with repositories, managing users and other facilities. As well as these important elements, *SCM* has an impressive branching system. Projects can be separated into different development branches and as such, development can be directed in different directions with the relevant branch.

Generally, source code management is done at the file level, and as such you need to deal with directory structures. *SCM* takes a logical approach to the user interface and has three separate panes for the branches, repositories and files. Dealing with checking in and out files and handling histories is good, but there were a few quirks to the interface. An example is that creating a branch took me a while to find out how to do it. I think Seapine may benefit from



Seapine SCM provides an intuitive user interface to source code management.

running the app by some usability analysts to help resolve these niggles.

One other small limitation I found in the software was regarding the built-in file editor – it is somewhat limited and offers no syntax highlighting and other editor features. This problem is only a small issue and a custom editor can be used for loading files into. It would be nice to incorporate some additional file viewing/editing facilities though to ensure the software can be used to quickly deal with editing files to fix small issues without the need of loading an entirely new editor.

I found all major tasks of source management (creating modules, adding code, committing changes, dealing with histories) easy to do, and the use of context menus across the application greatly simplified the learning curve of knowing how to conduct these tasks.

Seapine SCM & CVS

There is no doubt that the nearest competitor to *Seapine SCM* is the *Concurrent Versioning System (CVS)* that is freely available. CVS and *Seapine SCM* offer a similar set of functionality such as checking in/out files, branches, multiple users etc, although *Seapine SCM* differs in one major way – the interface. CVS is a command line tool, and although there are GUI clients for it, many are difficult to use (such as the venerable *WinCVS*)

and as such, many people resort to the command line tool. *Seapine SCM* alleviates this need for a good graphical client for source code management. Not only does *Seapine SCM* help in this area, but the general configuration of the source management server (Seapine's own server in this case) is much easier and entirely graphically driven also. This is an area where CVS has also proven difficult – it's configuration requires editing text config files. For most developers who are not incredibly familiar with source code management systems, a GUI client such as *Seapine SCM* will prove helpful.

Documentation

Documentation with *Seapine SCM* comes in the form of a number of manuals. The documentation is well-written and easy to use, and is more detailed than I was expecting. Another nice feature about the documentation is that there is some documentation available on the CDROM as a .pdf file.

Although the printed docs were good, I had difficulty getting the online documentation working. Even after some fiddling around with configuration settings, unfortunately I was unable to resolve this problem.

Seapine SCM is overall a well-written application. The product has a good selection of features, and will perform

most of the tasks that are required for source control management in a multiple developer environment. The interface of the client is well constructed and laid out, aside from a few minor problems here and there.

Conclusion

If you are looking for a solid source code management system, and don't want to rely on the command line nature of CVS, *Seapine SCM* is a capable solution. Not only will the software provide the functionality you need, but it's multi-OS (Linux, Windows, Mac OS) availability will ensure that developers scattered across different platforms can use the same source management system. The newer version 2 of *Seapine SCM* will be released in August 2003. [LXF](http://www.linuxformat.co.uk)

VERDICT

Features	7/10
Performance	8/10
Ease of use	8/10
Value for money	7/10

Seapine SCM is an easy-to-use and powerful source code management application that will work well in a multiple-developer environment. Recommended.

LINUX FORMAT RATING
8/10

JAVA APPLICATION DEVELOPMENT

JBuilder 9 Enterprise Edition



Paul Hudson reviews the Mercedes-Benz of Java rapid application development...

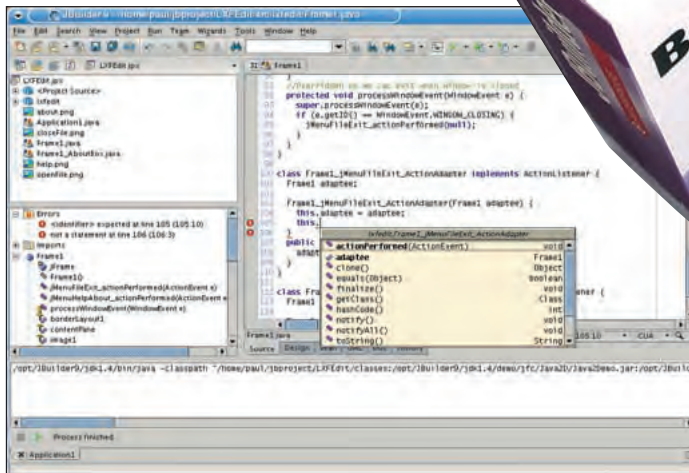
BUYER INFO

Rapid Java development environment.
See also *NetBeans*, *Eclipse*

- **DEVELOPER** Borland
- **PRICE** New user licenses from \$3,500, upgrades from \$2,990
- **WEB** www.borland.com/jbuilder/

How good is *JBuilder*? If you've never used it before, you've got a long way to catch up, because for a long time now *JBuilder* has been the definitive integrated development environment for Java, not least because of its advanced form designing and code generation functionality. As such, here at *LXF Towers* we're always eager to see what wonders Borland has managed to put in to the software to justify each version hike – we have yet to be disappointed, and happily version 9 is no different.

If you're totally new to *JBuilder*, let us bring you up to speed – drag-and-drop GUI creation, code refactoring, automatic error detection and fix suggestion, built in code optimiser and profiler, support for *Ant* and *Struts*, code completion and syntax highlighting, as well as built-in class-file obfuscation to stop reverse engineering. *JBuilder* has always been a very complete package designed by the



Wait, I got all this and haven't typed anything yet?

undisputed leaders of RAD development, and version 9 only pushes their lead even further.

In version 9, a host of new features have been implemented to improve the system across the board, including full integration of Borland's *Optimizelt Suite* product that helps you write faster code, better smart code templates (including macro support), enhanced error checking and correction, as well as tweaked CVS support.

Perhaps the coolest feature in version 9 is called 'Sync Edit', which offers the ability to edit a piece of code, and have *JBuilder* automatically update

all matching pieces of code – in effect, you can change a hundred instances of a chunk of code just by editing one.

In addition, *JBuilder 9* runs entirely on Java 1.4.1, so is fully compatible with the latest additions to the language. Being Java-coded, it looks and works the same irrespective of the platform on which you install it, and the Linux port works perfectly.

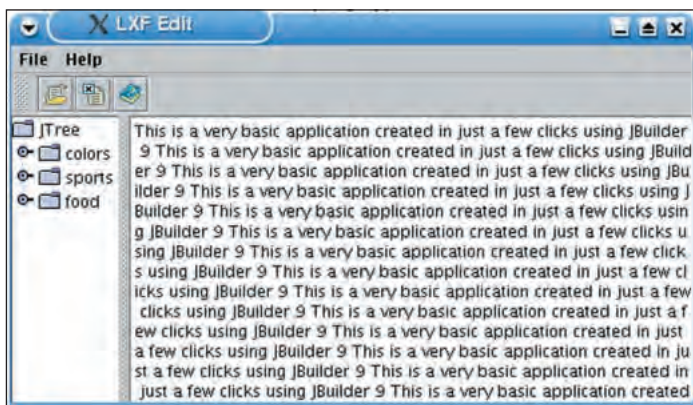
Menu creep

Using *JBuilder* has always been an enjoyable experience for me, and this version continues the tradition nicely. Performance continues to be just a touch too slow for the application to feel properly responsive as you'd expect on today's computers, but it is still better than previous versions – Borland seems to be continuing to work on tweaking the performance wherever it can. One minor problem that is common to apps that are as mature as *JBuilder* is that there is 'menu creep' – the process of menu bars getting longer and longer without any major redesign taking place. *JBuilder 9 Enterprise* now has no fewer than 25 options under the 'Tools' menu, for example – far too many to be user-friendly, and really in dire need of some

sorting out. With the ability to create native executables, CORBA interfaces and applications, as well as increasingly powerful web services, *JBuilder* continues to have the lead when it comes to features – the fact that each of these are available for creation though a friendly and simple GUI is merely icing on the cake. If you're new to Java, there's still lots to be had – *JBuilder* comes with literally dozens of tutorials to guide you from a basic user's point of view in how to get to grips with the most complex features of the package.

In the box, despite this being the most expensive edition of *JBuilder*, there aren't any printed manuals – this is quite sad, really, when one considers that *Kylix* comes with such excellent documentation. However, this is mitigated somewhat because *JBuilder* does come with excellent online help covering all parts of the product in very comprehensive detail.

Overall, this is another strong upgrade for the *JBuilder* line that's well worth the cost if you're a serious Java developer. Borland is pushing team-based programming harder than ever, and the end result, when combined with the new improved error detection, is that you can turn around projects faster than ever. [LXF](http://www.linuxformat.co.uk)



A quick app created using *JBuilder* just by drag-and-drop through the GUI.

VERDICT

Features	10/10
Performance	8/10
Ease of use	10/10
Value for money	10/10

Every bit as good as its predecessor, with some excellent additions to boot.

LINUX FORMAT RATING
9/10

BUSINESS LINUX DISTRO

SuSE Enterprise Desktop 1

Can SuSE's new business-centred distribution set the corporate world on fire, or are the developers perhaps spreading themselves too thinly?

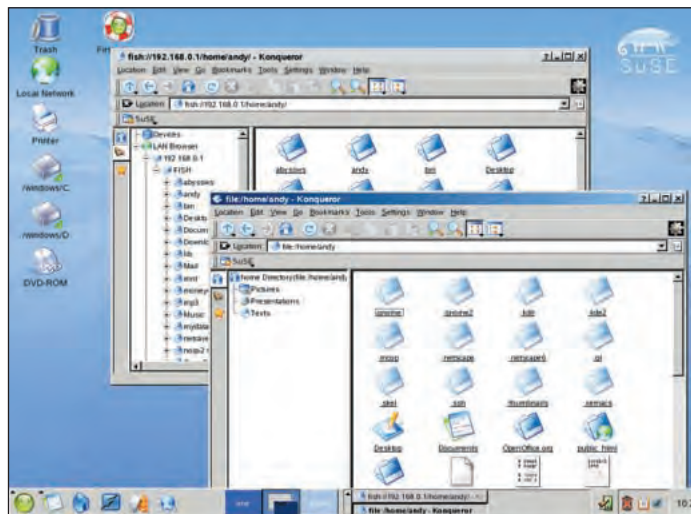
Andy Channelle takes a spin round the floor with the desktop debutante.

BUYER INFO

Linux distribution designed to make life easier for both administrators and users migrating from Windows and MS Office.

- **PUBLISHER** SuSE
- **PRICE** £410.00 (£481.75 inc VAT) for five machines
- **WEB** <http://www.suse.co.uk>

While SuSE's Personal and Professional product lines head toward double figures in the version number stakes, the company has decided to hit the drawing board to create a new boxed product aimed specifically at enterprise buyers investigating alternatives to Microsoft Windows. And while SuSE Linux Desktop has many similarities with the most recent standard release, it has been given the version 1 designation to set it apart from the its progenitors. It will, SuSE says, have a longer release cycle, more enterprise-friendly support and maintenance options and an application bias toward the standard office jobs. Moreover, the package includes tools developed to make rolling out SuSE desktops across hundreds, or thousands, of machines



The LAN browser allows you to view and access the entire network.

simpler while centralising many of the routine setup, configuration and general administration jobs. Finally, to make adoption smoother, it also includes a number of commercial applications – Sun's *StarOffice*, CodeWeavers' *CrossOver Office 2.0* and a small but useful selection of professional typefaces from AGFA | Monotype – which should make migrating both user skills and legacy documents less of a chore. It is arguable that SuSE may get more

mileage from including something like *Ximian Connector* (which allows *Evolution* to work as a client for *MS Exchange* server) in place of *StarOffice*.

As usual, SuSE has done a very good job of presentation, the box is packed with five CDs, a pair of manuals helpfully split along 'user' and 'admin' lines and, in lieu of the traditional stickers, a rather nice branded SuSE mouse mat.

Installation

The first real difference between this and SuSE's established distributions – apart from the fact that this is dropped back to kernel 2.4.19 – comes in the installation classes on offer. Firstly you get three options:

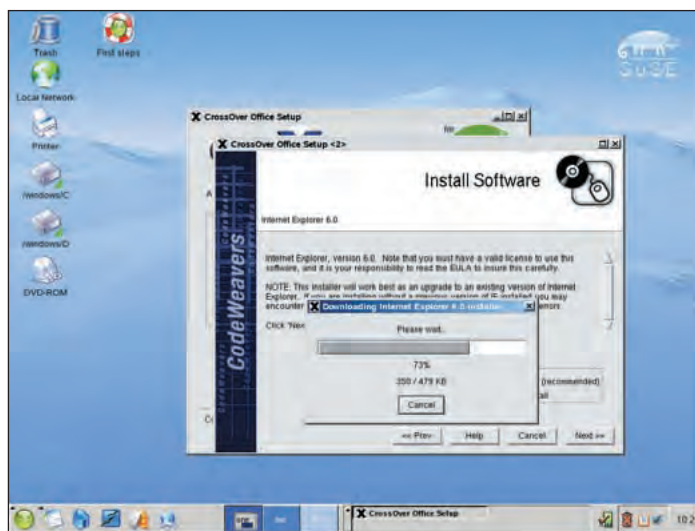
Single PC or small to medium-sized Office Network – SuSE says this is ideal for standalone machines or offices with a network of up to 50 machines. As applications and data are stored locally, you'll need about 1.4GB of free space for this option.

Template for Enterprise Clients – This class is specially designed for deployment on large networks installing, still locally, a reduced selection of basic applications including office package, email/web client and SAP client.

Template for Thin Client or Slower Computers – This option includes a small-footprint desktop and has been developed for situations where applications and data will be largely run from a central server. It should consume about 500MB of space.

Once this is passed you can choose to run either KDE (the standard SuSE desktop), KDE optimised for Windows users or GNOME. After fiddling with various selections to see the scope of them, I selected *Template for Enterprise Client* and *KDE for Windows Users* with the target audience in mind.

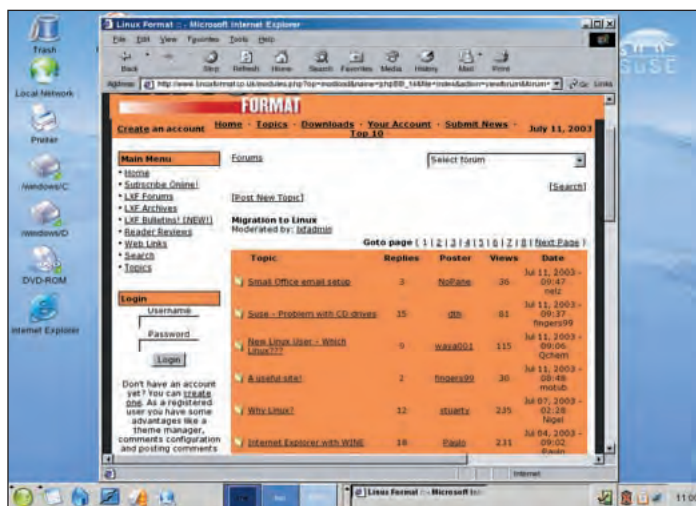
Installation follows the same routine as SuSE 8.2, meaning very good hardware detection, intuitive partitioning options and lots of feedback. During the X setup phase on a laptop install, the system ground to a halt after a 'successful' resolution/colour depth test which needed a hard reboot to solve. It didn't seem to affect *YaST*'s flow, though and the rest of the install passed without incident. The problem wasn't reproduced on any other test machines and a second installation (with the same configuration options) on the same laptop went smoothly.



CrossOver Office 2.0 adds Windows compatibility for many applications.

Support SuSE helps out

The SuSE Desktop boxed product includes licenses for five machines. This includes licenses for the extra commercial apps, a year of maintenance and client licenses for connection to SuSE's Open Exchange server. The maintenance agreement covers bug fixes, online updates and access to a special website. Further maintenance contracts can be purchased on an annual basis with prices starting at £335 (£393.63 inc VAT) for up to five clients. While buyers of SuSE Linux Personal/Professional are limited to 60/90 days of installation support, SuSE Desktop has this for the life of the agreement.



Internet Explorer proved to be a little flakey, but many other Windows applications ran without a problem when we tested them.

Input was needed only for specifying network information – IP address, netmask, gateway etc – and when *YaST* tried to foist a 800 x 600 screen resolution on me.

Ease of use

Like the manual, SuSE Desktop has been developed with two distinctly different groups in mind: users and system administrators. The desktop is designed to be familiar to those trained on Windows, and as such is launched with 'Redmond' window decoration and the QT Windows window style. This ethos is also clearly reflected in the start menu structure with its Applications, Settings and Documents entries. In our unscientific tests, Linux virgins could find the file manager, word processor, email and web browser without struggling, and managed to open and print mildly complex *Word*, *Powerpoint* and *Excel* files without any help.

As well as the regular Applications menu, there is also a Windows Applications entry where programs installed under the bundled *CrossOver Office* are stored. Once you've installed applications in here – *CrossOver Office* supports *MS Office 97, 2000* and, to an extent, *XP*, as well as *Internet Explorer*, *Adobe Photoshop*, *Lotus Notes* and *Quicken* – they follow the menu structure of Windows. Though most things worked as expected, *Internet Explorer 6* was quite unstable (this wasn't evident in a standard *CrossOver Office* installation) and sometimes required the use of **xkill** to clear the screen.

Browsing is handled by either *Konqueror* or *Mozilla*, with the former given prominence on the taskbar, while *KMail* takes care of email chores. I don't know if *Mozilla* was included for compatibility reasons but I think, with the noble intention of

making Linux foolproof, novice users could probably do with less choice.

And so it is with office suites. *OpenOffice* is installed by default, which seemed a little odd as one of the big features of the product is a *StarOffice* license. But there is a problem with *StarOffice*; not a giant one, but it is very annoying and adds time to configuration. I installed the suite via *YaST*, but when I launched the word processor, I was bounced into the *StarOffice* setup script, which did another install in my /home directory.

And then the real snag hit: as far as the system is concerned, *StarOffice* is installed in /etc/opt which is not writable; so, after typing a few hundred words, I attempted to save, causing the application to crash taking my precious work with it. Removing the suite via *YaST* leaves it installed in /home, though to access it you need to rebuild the menus. Like I said, annoying and time-consuming. Also running *StarOffice* next to *OOo*, in all its anti-aliased finery, demonstrates how far the Open Source version has developed.

The users' manual is well written and targeted, covering everything from the basics of the KDE desktop to more advanced tasks such as file encryption (which this package makes a breeze, by the way).

Administration ABC

While attempting to make Linux suitable for the average Windows user, this package also bulks up the tools available to sysadmins to make sure that they:

- A** can create a unified environment without configuring a thousand machines individually; and
- B** limit the scope of damage capable of being wrought by the average Windows user; and
- C** sort out any problems that arise without having to actually venture out onto the shop floor. Though, to give credit where it is due, much of this functionality comes courtesy of the KDE development team and the *Kiosk* project in particular.

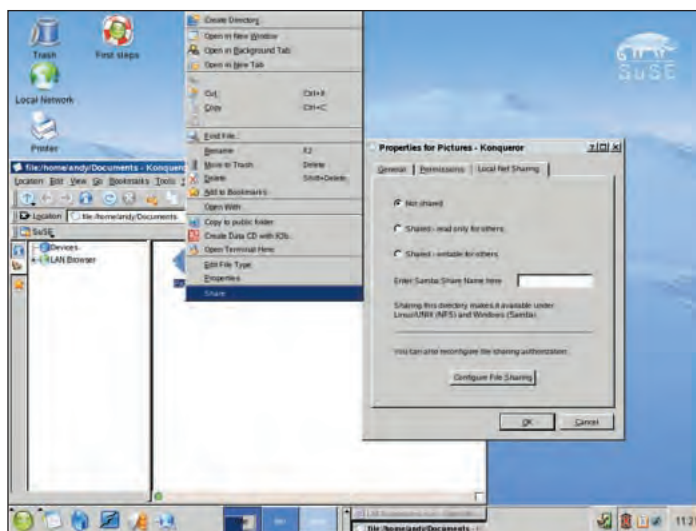
The first problem is solved by KDE's excellent central configuration options: the administrator configures a standard desktop, defining menu structure, look and feel etc, and then moves the configuration files (from /etc/opt/kde3) onto an NFS or *Samba* share on the server. It's then just a case, when installing the client systems, of making sure this directory is mounted via NFS/*Samba* as /etc/opt/kde3.

Once the default desktop is set up, *Kiosk* allows you to then set limits on what users can and cannot do; eg it is possible to prevent the launching of apps not on the desktop, deny shell access (though this isn't secure) or remove the ability to log out or lock the screen. URL manipulations provide options for redirecting HTML requests to specific hosts to local files, or restrict a user to browsing their own /home directory. Currently *Kiosk* is configured using a plain text file (*kdeglobals*), but the commands aren't too obscure and the Admin manual gives a good account of what you can and cannot do.

If things go wrong, or networked machines need some work, you can get access to the entire network through the LAN browser. This uses the Lan Information Server (LISa) daemon to periodically scan the network and list all connected devices (it does generate a lot of network traffic though). Access is then provided via *fish* which gives secure and authenticated access to the machine as either user or root. The use of the LAN browser is limited somewhat by its use of IP addresses instead of host names to identify network elements; this is OK on a tiny network, but it doesn't scale well.

Some networking additions do work well. The directory/file share options follow the Windows XP path of allowing or denying shares from a well-designed dialog box, letting the user decide whether the shared file is read only or writable and making accessible by either NFS or *Samba* with just a few clicks. The Remote Desktop Sharing feature is also well implemented. These are genuinely useful additions.

Finally, clients can be set to be automatically updated over the LAN, ftp or http, making local testing followed by a network rollout of updates efficient and straightforward to perform. **LXF**



Sharing files and directories has been simplified a great deal.

VERDICT

Features	7/10
Performance	6/10
Ease of use	8/10
Value for money	8/10

Issues with *StarOffice* mar what is otherwise a very good introduction to Linux for corporate buyers. It is an excellent showcase for the KDE's new *Kiosk* features though.

LINUX FORMAT RATING
7/10

ReviewsXimianDesktop2

DESKTOP DISTRO

Ximian Desktop 2

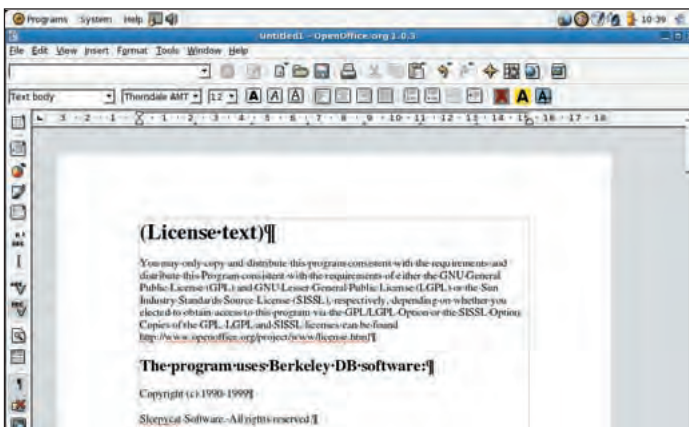
Monkey see, monkey download. Richard Cobbett does...

BUYER INFO

An entirely new desktop with the goal of making Linux attractive to the office, without taking away the power we all know and love.

- **SUPPLIER** Ximian
- **PRICE** \$99
- **WEB** www.ximian.com

Without a doubt, the best part about using a Linux desktop is that you can do almost anything. Want your *Kmenu* to stick out of that corner? No problem. A main menu? Easy. Reproduce the visual stylings of Windows or Mac OS? You could do that too, if you really felt the need. A tweak here, a new icon set there, the right background and you're finished. Unfortunately, in this one instance, freedom has a tendency to put a bullet through practicality. Even using your distro's vanilla settings, it's not uncommon to see clashing icons, garish colours and the feel that your desktop



OpenOffice.org may not look much different here, but compare it to the regular downloadable version and you'll soon see what an overhaul it is.

has been hastily welded together rather than lovingly crafted. Ximian Desktop attempts to fix this, merging all of your favourite Linux programs into one smooth, consistent, and above all else professional-looking interface that would be as comfortable on your office desktop as it would on your home PC. It's not

simply the Linux desktop that you can take home to show Mother, it's the one you'd be happy to present to your boss.

If you've tried installing a new desktop from scratch, you'll know what a painful process it can be. Ximian isn't completely trouble-free, but they're patience related rather than technical

annoyances, and overcome with time rather than textbooks. The easiest way to install the latest version is online, requiring you to simply switch to root and type **wget -q -O - <http://go.ximian.com> |sh**, but be warned that you'll be in for a very long wait. Ximian is installed using the *Red Carpet* package manager, and it took three attempts to pull down all of the required files – including many heavy hitters, such as *OpenOffice.org* and the GNOME environment. While this is obviously no problem on broadband, and a successful installation can be done in under an hour, you really don't want to try it with a 56K modem.

Auto configuration

Once the files are downloaded, Ximian automatically alters your configuration files, presenting you with a stylish new logon screen the next time that you boot up and calmly ushering you onto your new desktop. We installed it like this under both Red Hat 9 and SuSE 8.2 Professional, and it couldn't have been easier, but there is also an option to install from CD. A full copy of Ximian Desktop costs \$99, and comes with a handful of useful, if basic extras, such as copies of the Flash player, Windows-compatible AGFA fonts, a year's subscription to *Red Carpet Express* and the customary 30 days-worth of support.

Once installed, Ximian is striking – both in terms of aesthetics, and its difference from any other desktop around. Most notably, it feels as though the screen has been flipped over, with your program menu and shortcut buttons hanging from the top of the screen, leaving the bottom free for the taskbar and virtual desktop switcher. Ximian is ultimately GNOME remixed, so you can certainly dive in, add your favourite panel applets and generally shunt everything around, but you're intended to use the system more-or-less as provided. Certainly, it's only by doing this that you'll get the full benefit out of it. Each desktop nip and tuck may be small, but they work together beautifully – from the consistent icons on your application menu to the



Ximian Desktop's stylish windows theming ensures that every program fits on the screen, whether you're playing a game or sending bizarre messages.



The integration between Ximian Desktop's applications ensures that everything you do simply works, without technical messing around.

unobtrusive – but never fiddly – shortcuts at the top right of the screen. By default, these are created for the web-browser *Galeon*, *OpenOffice.org* and *Evolution* – Ximian's wonderful *Microsoft Outlook* replacement – but others can be added simply by dragging and dropping icons from the menus.

These menus too have seen a hefty polishing job, using larger than normal icons, and splitting oversized lists with a 'More' option to prevent your screen collapsing in on itself after one installation too many. Those installed by default are principally listed by their function rather than name – *Music Player* instead of *XMMS* and *Instant Messenger* instead of *GAIM* – making it much faster to get up and running.

Naturally, this means that one application of each type is provided, but alternatives are only a quick click away. The *Red Carpet* application that you use to install Ximian Desktop in the first place puts in a second appearance as 'Software Update', working behind the scenes to download the latest patches and software for both your distribution and Desktop itself. Operating like the traditional *Synaptic/apt-get* combination, *Red Carpet* automatically handles the more painful aspects of Linux software installation – from downloading the correct files for your distribution to fixing troublesome dependency issues. The only thing that it doesn't do is automatically add the installed application to the appropriate program menu for ease of access, which can be troublesome if you don't know the name of the command that you need to run. A premium version of this service, *Red Carpet Express*, is sold separately, and offers high-bandwidth servers to pull your downloads from, but this free version works perfectly well

– scoring under five minutes on the internationally recognised *Time To Frozen Bubble (TFB)* scale.

Configuration has seen a similar overhaul. All system controls are now packaged on two screens, one for regular users to play with (controlling fonts, background images, accessibility controls and other such tweaks) and another for system administrators (SuSE 8.2 vintage *YaST* in our case). Network settings depend on the distribution that you're using, but anybody can take advantage of the new Printer Wizard, which takes you step by step through the process of finding and configuring local and network printers alike.

The most exciting part of Ximian Desktop however is its redesigned version of *OpenOffice.org*. One of the biggest subconscious draws of Windows packages is that they automatically fit. *Microsoft Office* never feels like an interloper – everything from its icons to menus seamlessly clicking with its surroundings. Done correctly, you don't even notice it happening, but its absence sticks out like a sore thumb. By default, *OpenOffice.org* feels as if

someone's smashed down a rather weighty sledgehammer on the entire hand, but this version is rather special. By focusing specifically on Desktop, Ximian has been able to delve into the deep into the various applications and completely integrate the world's best Open Source office suite. Over 800 shiny new, alpha-blended icons are provided to match Ximian's aesthetics, but its beauty is not simply skin deep. Click on an email address, and up pops an *Evolution* message composer. Choose a URL and *Galeon* puts in an appearance to take you straight to the site. Most controversially, it comes configured to save documents in *Microsoft Office* format. While this initially seems to be letting the side down, it makes perfect sense in practice – Ximian is an office desktop, and setting *Microsoft* to be the default guarantees that anybody will be able to read the files that you send out. *OpenOffice.org* is good enough to handle almost anything sent back, enabling you to focus on your actual project rather than the technical headaches of cross-platform communication.

Integrated printing


Back on your own computer, *OpenOffice.org* boasts full CUPS integration for printing, and the ability to access remote files as easily as selecting them from your network, remote directories and windows shares. The combined result is that it's no longer a simple application, but an integrated part of the operating system – exactly the kind of tweak necessary to compete in an office environment. Just to put this into context, an installation of Ximian Desktop gives you a full *Microsoft* compatible office suite, *Outlook*-class email client (with a

separate *Connector* plugin that enables it to link up with full *Exchange* servers), a web browser capable of taking almost any site we threw at it, and an idiot-proof method of installing new software – and all this straight out of the proverbial box.

What downsides?

The only real downsides to using Ximian Desktop are those of personal preference – most notably the choice to use GNOME rather than KDE. While both offer excellent desktops, the latter has long-been the most popular in the Linux world, and the absence of a few of our favourite features, such as easily setting up alternate backgrounds for each virtual desktop or keeping useful applications on the top of the screen, did prove irksome at times. Equally, for all of Ximian Desktop's refinements, you do occasionally find an application throwing a monkey-wrench into the works, most notably *XMMS* and its rabid refusal to click to the bottom taskbar, instead of where a 'normal' sized KDE panel would have ended.

Do these sound like trivial complaints at best? Absolutely – and perhaps the best compliment that we can pay to Ximian is that such quibbles proved the only real low points of our extended use. Everybody will have their own handful, from the technical method that Ximian used to dress up *OpenOffice.org* to the fact that you can't rename or Trash the 'My Computer' icon on the desktop, but there is little that will slow you down in the long term. Ximian's GNOME base means that while this is only the second big release, its individual components have been tried and tested. The fixed system that it brings to the table probably won't do much for you if you insist on retaining total control of your system, but there's no better way to give your machine the professional makeover that it deserves. **LXF**



Red Carpet is on hand at all times to download and install new software. There's been reports of it clashing with *apt-get*, but it's still a useful utility.

VERDICT

Features	9/10
Performance	8/10
Ease of use	8/10
Value for money	9/10

Linux has never looked so good, but you may prefer a fraction more control over your home box.

LINUX FORMAT RATING
9/10



SERVER APPLIANCE

Navaho TeamCAT

Looking for a Linux appliance that handles web email, spam and virus filtering, firewalling and VPN, and does so while looking great? **Paul Hudson** has found the solution you're looking for...

BUYER INFO

Powerful network management device for 10-500 clients.

- **DEVELOPER** Navaho
- **PRICE** from £2030 (business)
Education pricing scale also available
- **WEB** www.navaho.co.uk

Navaho has been producing hardware solutions for quite some time, selling mainly into education and SMEs.

Their CAT series, which stands for Complete Appliance Technology, comes in three flavours, NetCAT, MailCAT, and TeamCAT, of which the latter is the most expensive and the one we received for review. Each CAT model offers its own level of software capability, culminating in the TeamCAT which is practically a plug-and-play "do everything" box, including working as a fax server, bandwidth manager, and groupware server. Each machine is fully standalone, but Navaho has built failover backup into them so that you can plug two together that keep themselves synchronised together – as soon as the "heartbeat" signal from the master fails, the backup takes over automatically.

Inside the machine is powered by a fast P3 and a customised Red Hat 2.4.18 kernel. Navaho has done a lot of work on customising the kernel, and have even had patches accepted into

the stock kernel itself. A lot of the software they have running is off-the-shelf GPL software that they pre-configure to make it easy to work with no fuss. For example, *iptables* backs the firewall solution – Navaho has it set up to block all ports but the important ones, and offer an easy-to-use interface to manipulate the policies. The server is also set up to automatically audit and log all parts of the system so that administrators can keep tabs on its operations with a few clicks.

The big feature of the Navaho solution is the technical support – it's a comprehensive package spanning from automatic vulnerability monitoring and software updating to extensive help desk support. They have two levels of support, standard and managed, and both include remote software updates – as new versions of software you run come out, you can configure your box to automatically email the administrator detailing what's new in the package. Alternatively, as part of the managed support, you simply give control of the box over to Navaho, who will upgrade and patch software as appropriate. The anti-virus definitions are also updated remotely, with the box synchronising with Navaho every hour to make sure it has the best possible defence.

For tech support queries, Navaho offers support through the web, by email, and by telephone, with the

telephone option restricted to 8 to 6, five days a week. The engineers there are qualified to deal with all varieties of questions involving the unit, and can even remotely fix problems if the on-site administrator grants them access. This is very impressive given that the quality of tech support in other companies seems to be continually declining – Navaho is taking a strong lead, and offering a service that means you can actually rely on its support staff to the point of thinking of them as outsourced server admins.

Won't cost the earth

The charges for TeamCAT are very reasonable, at £2030+VAT for a unit licensed to handle 1-10 workstations, including technical support. The cost rises to £18,310+VAT for a unit licensed to handle a maximum of 500 clients, which is the maximum the hardware is capable of handling. A lot of the price is because of the technical support – naturally a company with 500 clients will generate more support calls than one with just 10 clients. However, even at the highest level you're still only paying around £36 per client per year, which is excellent. You can purchase a second TeamCAT to set up the backup cluster for just another £1500 irrespective of the number of clients connecting.

Once your initial technical support runs out, you pay yearly for support.

For 1-10 clients on the standard support package, you pay £555, and that rises to £9,100 for 500 clients. For managed support – where Navaho will manage every aspect of your server for you – the costs start at £1000 and rise to £16,380. Given that Navaho's support is so very good and that the box has so much functionality packed into it, these prices are something of a bargain.

Overall, the Navaho machine is a powerful little unit, cramming email delivery, web caching and filtering, and security features into just one small box that literally runs itself. If you want even less to do with the machine, Navaho will even manage it for you remotely eliminating all doubts and worries. If you're looking for a box that can revolutionise the way your company works *without* a hefty price tag, look no further than this LXF *Top Stuff* award winner. **LXF**

VERDICT

Features	10/10
Performance	9/10
Ease of use	10/10
Value for money	10/10

An impressive unit that crams a lot of work into such a small form factor.

LINUX FORMAT RATING
 **10/10**

Programming C#, 3rd Edition

Paul Hudson has been idly wiling away the long summer afternoons reading the latest addition to O'Reilly's .NET series...


BUYER INFO

■ **AUTHOR** Jesse Liberty
 ■ **PUBLISHER** O'Reilly
 ■ **ISBN** 0-596-00489-3
 ■ **PRICE** £31.95 **PAGES** 689

As languages go, C# has received an awful lot of criticism for being too close a copy of Java. However, it's easy to argue that this state of affairs is largely unjust – Java itself was of course a fairly close copy of C++, which was in turn an improvement over C, so it's hardly a bad move to base C# upon the errors and weaknesses that have to come to light in Java since it was launched.

Jesse Liberty, the long-standing book author and language expert, brings his full weight to bear in the third edition of this very popular book, and I'm happy to report that it has lots to offer newcomers to the language as well as owners of previous editions of the book. Liberty takes the smart approach of teaching the C# language in conjunction with the .NET class libraries, which means that readers learn how to write C# alongside how to use their new-found knowledge can be used to get powerful results. With the author bringing such a huge amount of writing and technical experience to the arena, this book can't help but be authoritative to its core – I don't think I've ever read one of Jesse Liberty's books and found what he says to be wrong!

Because the book starts at a very low level, it's a safe bet for programming beginners and those who have just a little experience. From such lowly beginnings, though, the book manages to reach great heights – complicated topics, including such things as operator overloading, regular expressions, database access, and data marshalling, are all covered and conquered as you turn from page to page. While there is no real content for very advanced programmers (a section on code optimisation would be most welcome), there is a full description of the language from absolute beginners level all the way up to its most complex features. Given the



relatively short length of this book, the depth of knowledge covered pays tribute to Liberty's "to-the-point" writing style.

One aspect I particularly found useful were the insightful and numerous comments sprinkled around the book directed at C/C++, Java, and Visual Basic programmers, that offer helpful advice to make learning C# as simple as "This is how you did it in Java; this is how it's done in C#". As the author states near the beginning, many VB6 programmers will find it easier to switch to C# rather than VB.NET – it's good to see that the transition from other languages has been made so smooth. One of the most interesting parts of the book is its coverage of web services using C# and ASP.NET – if you've ever wanted to learn more about web services from a very practical perspective, this book is certainly a good buy.

The tone taken through the book is quite chatty, matching Liberty's usual style, and surprisingly you find you've learned a lot more than you thought was possible – the friendly tone doesn't water the content down at all (as is does in so many books in the computing sphere aimed at novices), meaning that despite being under 700 pages you manage to create some fairly complicated applications with the code presented.

Being the third edition, there are few if any errors in the code samples that are presented throughout the book, although noticeably obvious is the fact that *Programming C#* doesn't come with a CD so you'll need to

enter the code in yourself by hand or download it from the web – downloading is definitely the preferable option, as some of the code blocks are many pages long! When I went to download the source code, I sadly got a 403.9 error message back from a Microsoft IIS web server, saying that there were too many people accessing the site – very frustrating, particularly considering the lack of CD, and also a little ironic given that the book's focus! At least it shows that it's popular...

The biggest drawback to the book is that it is very firmly focused on Windows developers, despite Mono's increasing popularity – there are references to Microsoft's Visual Studio .NET programming environment throughout, with all web screenshots taken of Internet Explorer. Mono is in fact not mentioned at all, despite much of the C# and ASP.NET code working perfectly well – this is a great shame in what is otherwise a fairly comprehensive book. At the very least, it would have been helpful to mark what works and what doesn't, so that people have some choice.

If you bought one of the previous editions, you're not short of good reasons to upgrade – not only have all examples and text references been updated to version 1.1 of the .NET framework, but large sections of the text

have been rewritten and improved. Particularly of interest will be the new chapter on delegates and events – this is much more detailed and easy to digest than the text from previous editions. One minor nag that still has yet to be addressed is the addition of a more comprehensive index – sixteen pages of index for a 700-page book is a little weak, and makes hunting around for a particular topic a bit slow.

Overall, *Programming C# 3rd Edition* is a solid improvement over its predecessor, and arguably the most complete and informative book on the topic available. It's not perfect, no, but it introduces the topic and takes you to a fairly advanced level at a good pace, and rarely fails to be enjoyable.

If the 4th Edition, whenever it's released, manages to at least mention the Open Source alternative – although more complete information would be better – that would definitely merit a Top Stuff award because finally the book would be complete. However, as things stand, this book falls short of greatness by a penguin's breadth – it's a solid reference guide that you'll keep returning to for years to come, but needs that little extra "oomph" to win us over completely.

VERDICT

An enjoyably comprehensive book on the topic, but you'll need to figure the Linux information out for yourself.

LINUX FORMAT RATING
 /////////////// 9/10

FreeBSD Unleashed, 2nd Edition

Linux isn't the *only* Open Source operating system available – Paul Hudson has been reading the new FreeBSD book from Sams...


BUYER INFO

- **AUTHOR** Michael Urban and Brian Tiemann
- **ISBN** 0-672-32456-3
- **PRICE** £36.50 **PAGES** 970

Despite Linux grabbing the lion's share of Open Source headlines, BSD has been popular for a long time.

FreeBSD, the most popular release of all the BSDs, has nearly as many programs as Linux (including KDE and GCC), as well as a much more mature codebase than Linux, and yet many Linux users barely give it a look. With this book, it's fair to say that "lack of documentation" is no longer a worry – at almost 1000 pages, this book covers FreeBSD from top to toe.

Sams' *Unleashed* series has an excellent reputation – if I were ever to write a book on PHP, I'd want it to be a part of this series! *FreeBSD Unleashed*



follows the same tried and true method – start simple, work smoothly up to expert level, then spend the rest of the book covering real world application of knowledge. As such, the book starts off discussing the pros and cons of FreeBSD vs Windows and Linux, as well as against other BSDs, then proceeds onto installation, basic shell manipulation, and then admin. By the time you're half way through the book, you've already learnt how to install

software, how to program using shell scripts and Perl, and how to optimise your system – no mean feat!

Towards the back of the book there's the real world information – a lot of text has been devoted, wisely, to how to set up a web server, a mail server, and a database server; how to configure NAT, a firewall, and a VPN server, with the book finally closing on how to full configure X-Windows. This is a *very* heavy read – not *boring*



heavy, but *information* heavy; there's very little space wasted, and you're taken on a rollercoaster of learning as you go from being an absolute newbie to someone who feels like they've been using the system for years.


If you're a Linux user wanting to try something different, or if you're wondering what people have been shouting about on Slashdot all these years, give this book a try – much of the information given will be familiar thanks to the shared heritage of the two OSes. If you're already a veteran BSD user, it's very likely this book will teach you many things you didn't know already – it certainly taught me.

VERDICT

If you buy only one BSD book, make it this one.

LINUX FORMAT RATING
 **10/10**

EASY TO BUY • EASY TO SET UP • EASY TO SEE



EasyVserver solutions
 Debian or RedHat O/S
 True "root" access
 4, 6 or 8GB raid space
 11P address
 Highly secure

p-p-p-pick up
 a dedicated linux server

Make the most of Linux technology with the big name in Web registration and hosting packages. Our flexible, scalable, secure EasyVserver solutions start at just £39 per month. Back up by unrivalled support and know-how. If you want the best of Linux come along for the ride.

log on today at: www.easyspace.com

Easyspace 
 your perfect partner for the web

Established in 1997 with over 1/2 million customers. Accredited ICANN registrar & nominet member. Prices exclude VAT.


More Java Pitfalls

Richard Cobbett reaches for a lifeline like Indiana Jones in the Temple of DOM.

BUYER INFO

■ **PUBLISHER** Wiley
 ■ **AUTHOR** Michael C. Daconta, Kevin Smith, Donald Avondolio, W. Clay Richardson
 ■ **ISBN** 0-4712-3751-5
 ■ **PRICE** £29.95 **PAGES** 480

More *Java Pitfalls* is, as the name would suggest, the sequel to the popular *Java Pitfalls*. If you're familiar with this first publication, little has changed. In the course of its pages, we see fifty potentially serious problems that you may encounter while programming professional Java applications, and full examples showing how to beat them. This rather more specialised than normal approach does have one obvious potential problem, namely that you may or may not hit these specific problems, but in working through them you will become more familiar



with the Java language and its many eccentricities, which could well have a similar effect on more practical problems.

The most obvious difference between the two books is that *More Java Pitfalls* focuses on the J2ME, J2SE and J2EE systems, while the original was all about .lang, .util, .io and GUI messes. It's also split into three distinct sections – the first concentrating on the Client Tier, with application development, DOMs, mouse button control and similar

issues, the second moving across to the web, servlets, JSP and form validation, and finally onto the Enterprise Tier, with session, entity, Enterprise Java Beans and JDO pitfalls examined and corrected.

Whichever section you happen to be in, the style of the book is very welcoming – a chatty, but professional voice, with even the occasional pun (“Cache. It’s money”), ‘let’s face it..’ anecdote, and a healthy selection of Dijkstra quotes to kick off each

section. When getting down to business, the crucial factor is that *Pitfalls* doesn't simply explain a problem, it demonstrates it with practical examples. You aren't expected to know every single command off the top of your head, but the code provided is not heavily commented – it says ‘Now we'll be doing this’, does so, and you're expected to be able to keep up. With the specialised nature of the problems on offer, this fixes *Pitfalls* firmly at intermediate level – but a superb way of really diving into Java's underbelly if you already have basic knowledge.

VERDICT

A potential lifesaver in the event of troubles, but potentially too focused for your personal needs.

LINUX FORMAT RATING

7/10

Python in a Nutshell


Richard Cobbett would have harsh words for the KP snacks company in such an eventuality.

BUYER INFO

■ **PUBLISHER** O'Reilly
 ■ **AUTHOR** Alex Martelli
 ■ **ISBN** 0-5960-0188-6
 ■ **PRICE** £24.95 **PAGES** 600

In a nutshell, *Python In A Nutshell* serves one primary goal: to act as an immediately accessible goal for the Python language. True, you can get most of the same core information that is presented within the covers of this volume online, but this will invariably be broken into multiple files, and in all likelihood lacking the examples or exact syntax description necessary to truly understand a command. What both methods do have in common is that they're not the best way of picking up Python itself.

There is a short introduction to the workings of the language at the start of the book, including advice on compiling the software that you'll need, tracking down suitable IDEs and



taking the first steps on both Linux and Windows, but get past this initial hand-holding and you're expected to know what you're doing. The section on Tokens, for instance, is summed up in one paragraph, as “Python breaks each logical line into a sequence of elementary lexical components,” with other sections largely identifying where Python breaks away from other languages – such as ‘In Python there are no declarations’. This section is a good way of brushing up your Python

knowledge though, especially if you've been learning from various different sources.

The meat of the book however is the library guide. This is split up into logical sections, kicking off with File and Text Operations, and running through Threads and Processes, Debugging, Server Side Modules and XML, pausing only to take a brief overview of the Tkinter GUI builder, Jython (for running Python on a JVM) and finally distributing your extensions

and programs alike using *distutils*. Each individual section is further broken into types and alphabetical order, with the full version of the command, a brief description, and short examples of their use – be it demonstrating what the output will resemble, or providing a quick piece of garbage collection code to plug into your source.

As with the rest of the book however, these demonstrations are intended to boost your existing Python knowledge, not teach it from scratch – there are few comments, and no full project-tutorials, which makes this a book for intermediate users. **LXF**

VERDICT

Excellent look-up guide, although you may prefer more electronic sources for your info.

LINUX FORMAT RATING

8/10

Roundup

Every month we compare tons of software, so you don't have to!

OUR SELECTION AT A GLANCE

- Kate
- Joe
- Vim
- jEdit
- Emacs
- Nano

Text editors

Paul Hudson dons his flame-proof suit and tackles one of the most controversial round-ups in the Linux world – text editors...

After “Hello World!” and a Celsius to Fahrenheit converter, most programmers produce a text editor to advance their skills, so it's no surprise that there are an abundance of text editors available. Some go for the “*Windows Notepad* replacement” market, others go for the “it does everything you can think of” market, and of course there are all sorts of shades of grey in between.

Originally on Unix systems we had *ed*, the standard text editor, was the primary editor. See the box *Ed: the standard text editor?* for more information on this venerable piece of

software. Nowadays, each user has their own list of requirements for a text editor: does it have syntax highlighting for code, does it spell-check, can it show line numbers, and more. As such, there's a great deal of loyalty to each text editor, with many users firmly believing their particular version is superior and getting embroiled in huge (and some would say, rather pointless) flame wars. Here we look at six of the most popular editors from all the thousands that exist, each with their own take on what a good text editor should be.

Text editors have come on a long way in the last couple of decades. In

Mr Nice, his autobiographical book about his adventures in marijuana smuggling, Howard Marks recounts that back in the early 1980s he was one of the World's first customers for a basic computerised text editor/ WYSIWYG word processor. He appreciated the benefits in terms of costs and efficiency of being able to edit text in such a way for his import/export businesses, both legal and otherwise. At the time, such technology was beyond the reach of most businesses, as the basic hardware and software setup started at £6000. You don't need the wealth of an infamous international drug-trafficker these days to edit your text though...

‘Nowadays, each user has their own list of requirements for a text editor, like spell-checking, syntax highlighting, line numbers...’

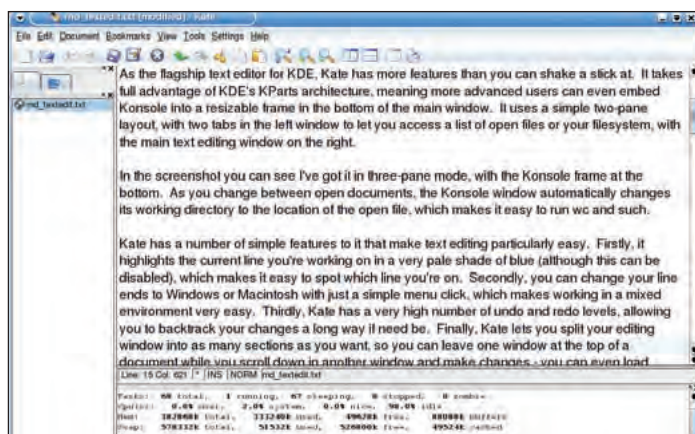
RoundupTextEditors

Kate

A text editor so powerful and friendly you may well want to kiss it...



■ VERSION 2.1 ■ WEB <http://kate.kde.org/>



As text editor interfaces go, this is top of the heap.

As the flagship text editor for KDE, *Kate* has more features than you can shake a stick at. It takes full advantage of KDE's KParts architecture, meaning more advanced users can even embed *Konsole* into a resizable frame in the bottom of the main window. It uses a simple two-pane layout, with two tabs in the left window to let you access a list of open files or your filesystem, with the main text editing window on the right. In the screenshot above, you

can see I've got it in three-pane mode, with the *Konsole* frame at the bottom. As you change between open documents, the *Konsole* window automatically changes its working directory to the location of the open file, which makes it easy to run the likes of *wc* and such.

Kate has a number of simple features to it that make text editing particularly easy. Firstly, it highlights the current line you're working on in a

very pale shade of blue (although this can be disabled), which makes it easy to spot which line you're on. Secondly, you can change your line ends to Windows or Macintosh with just a simple menu click, which makes working in a mixed environment very easy. Thirdly, *Kate* has a very high number of undo and redo levels, allowing you to backtrack your changes a long way if need be – something that users of all levels of ability will appreciate! Finally, *Kate* lets you split your editing window into as many sections as you want, so you can leave one window at the top of a document while you scroll down in another window and make changes – you can even load different documents into frames.

For programmers, there's a lot more to be had. For example, *Kate* has almost 60 different syntax highlighters, from C++ and PHP to Makefile and UnrealScript, meaning that it will almost certainly work well as a text editor for whatever language you're after. *Kate* also has line numbering and code folding for most languages, which, if you've never used it, allows you to "minimise" code blocks to just one line.

With all the features packed into *Kate*, it does take a little time to start up – almost a second on an 800MHz PIII, but that's hardly slow given all that it offers. For people who want maximum speed, a cut-down version

is offered in *KWrite* – it still has the main text window, but is missing quite a few of the extra frills, making it load a touch faster.

On the documentation front, *Kate* uses *KDE Help Center*, which is one of the parts of KDE that really lets the rest of the system down. By having help centralised, it's clumsy to navigate around, lacks search functionality, and is hard to use when you actually want to find information. Why anyone would want to read the Unix manual for *bash* when they click 'Help' from *Kate* is beyond us here at LXF.

Kate is missing a few little features that would make it perfect, for example it has no word count feature built in, which means there's no way to select some text and word count just that section – a minor pain. Other than that, it's a great program which is still advancing at a good rate, and deserves near top marks.

VERDICT

Features	10/10
Ease of use	10/10
Documentation	9/10
Performance	9/10

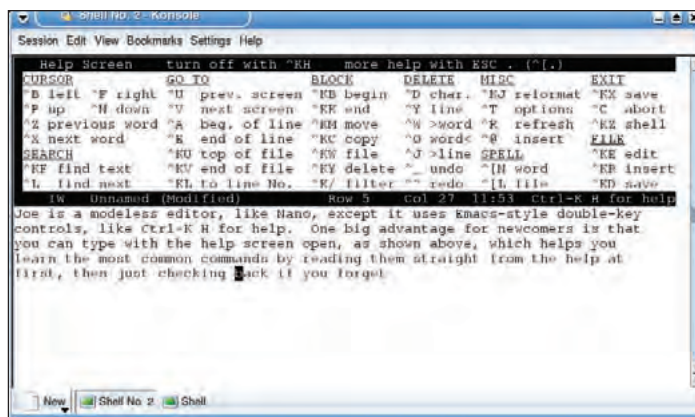
Feature-packed and easy to use, *Kate* is an excellent text editor for all uses that is still advancing in functionality.

LINUX FORMAT RATING
9/10

Joe

Joe's Own Editor is a mix of the others, with a popular following.

■ VERSION 2.8 ■ WEB <http://sourceforge.net/projects/joe-editor/>



Joe is a peculiar text editor

because it takes a number of ideas from *Nano*, a number of ideas from *Emacs*, then mixes them together with its own brand of usability. When we did a poll on the LXF web site about favourite text editors, we didn't include this one, and got justifiably flamed as a result – clearly it's popular!

Joe is a modeless editor, like *Nano*, except it uses *Emacs*-style double-key controls, like Ctrl-K H for help. Although this is complicated for newcomers, particularly when *Nano* makes shortcuts so easy, this is overcome to a degree because *Joe* allows you to type with the help screen open, as shown in the screenshot, which helps you learn the most common commands by reading them straight from the help at first, then just checking back if you

forget. Also, less dexterous users could be at a disadvantage trying to hold down all those keys!

Many people with more advanced usage in mind will probably find *Nano* a little too simplistic, in which case *Joe* is a great option for them – it's nowhere near as complex as *Emacs* or *Vim*, but adds just enough functionality to make the jump from *Nano* worthwhile.

VERDICT

Features	7/10
Ease of use	10/10
Documentation	8/10
Performance	9/10

A happy blend of the others, but lacking advanced features.

LINUX FORMAT RATING
8/10

It's popular with quite a few of LXF's readers, so here it is!

Vim

An improved version of *Vi*, but is it improved enough to come top?

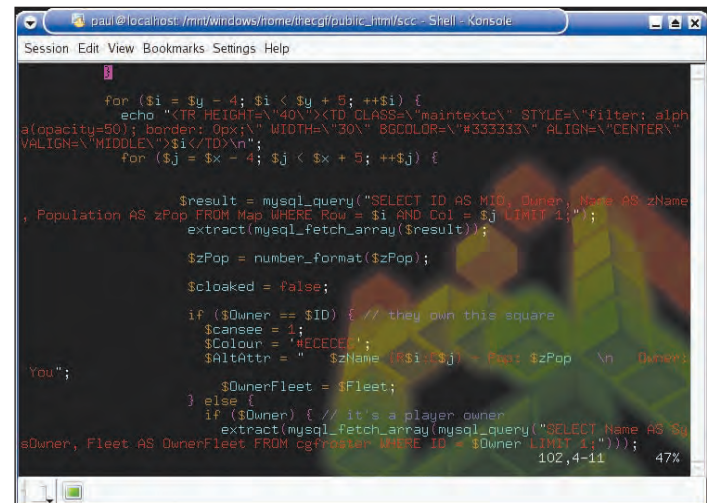
■ **VERSION** 6.2 ■ **WEB** www.vim.org/

Vim, an improvement over standard *vi*, is a powerful and yet fairly easy to learn text editor that is standard across nearly all Linux systems – no matter where you go, *vi* is almost certainly already there. Many newcomers find *Vim* hard to learn because of its very logical way of working – you mix numbers and single-letter mnemonics to get the actions you want, and there are such a variety of mnemonics available that many get lost fairly easily. For example, “dd” is “delete line”, “5dd” is “delete five lines”, etc.

When used as a basic text editor, *Vim* is fairly easy to learn and use. However, its method of copying and pasting is quite archaic, and it has a knack for doing entirely what you don't want it to do – recording a macro rather than quitting is the most popular! However, text editing with *Vim* is, on the whole, easy and powerful once users get the hang of switching between

editing modes and working with multiple lines at the same time. *Vim* is probably the most popular text-based text editor available, simply because the learning curve is quite low and yet you get a lot of the important power of *Emacs*. Granted, *Emacs* is so much more flexible that it can impersonate all of *Vim* smoothly, but the added complexity in the app often confuses people coming from *Vim*, and there is a fairly good *Emacs* emulator in *Vim* – *vimacs* – which keeps the playing field level.

Vim has good syntax highlighting for key languages, but often the colours are dark and hard to read. Line numbers, code folding, and other programming aids are present, but much of this is configured through files – there's no real easy way to configure *Vim* other than editing files directly. Code folding is particularly clumsy, and requires special syntax to be placed into the file. Two advantages to *Vim* are its availability and



documentation. *Vim*, or at least *Vi*, are available on most if not all Linux distros, so when you connect to another computer through SSH or telnet, you can be sure it'll be around. For docs, the *Vim Tutor* is a guide to all the key features of the program in one fell swoop, and usually gets people up to speed in no time – well worth a look.

Overall, *Vim* has always been a good editor to learn because it's available on most Unix boxes, and that's not going to change for a while yet. Indeed, *Vim* is undergoing a bit of a revival right now with the hope of a *Vim KPart* being in an upcoming release of KDE.

No matter how hard you try, *Vim's* syntax highlighting is never that easy on the eyes.

VERDICT

Features	8/10
Ease of use	9/10
Documentation	10/10
Performance	9/10

A stalwart of text editors in the Linux world, *Vim* continues to be as ubiquitous as ever.

LINUX FORMAT RATING
9/10

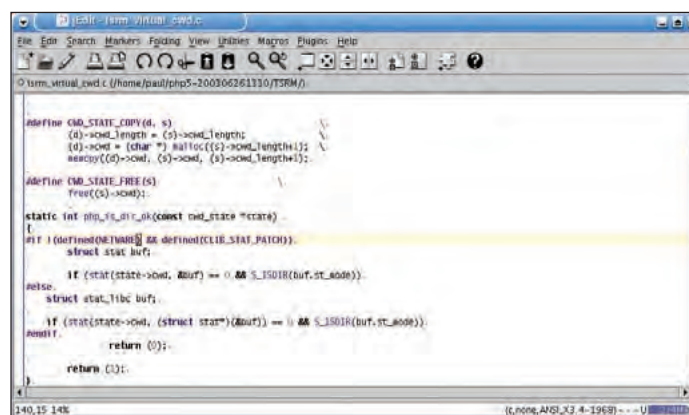
jEdit

The only program here developed in Java – does it give it any advantage?

■ **VERSION** 4.1 ■ **WEB** www.jedit.org

Designed primarily as a programmer's text editor, *jEdit* is Java-based with a raft of programmer-friendly features – large amounts of code folding variations, code markers, and search options, which add up to make it a powerful tool for anyone who does a lot of work programming. However, it doesn't have much in the way of general text editing functionality, and uses *Emacs*-style keystrokes to accomplish tasks. For example, to redo an action, you press Ctrl-E then Ctrl-Z – not what you'd expect, and not particularly easy to do for such a common operation. *JEdit* is missing a view of open files, although it compensates for this somewhat by having a drop-down box you can select an open file from.

The unnecessarily complex UI is a barrier for new users, and the Java user interface (regardless of skin choice)



A range of niggles forces *JEdit* to hide what light it does have under a bushel.

makes the editor feel alien no matter what platform you're used to. Furthermore, the horizontal scrollbar changes its size based upon the current lines on screen, which is bound to confuse people. On the plus side, *jEdit* has an excellent word count function (including selection word count), powerful macro support, and proper support for code indenting/unindenting, which is such a rare thing in these days. Furthermore, thanks to its Java nature,

jEdit is cross-platform and works smoothly on all Java-supported platforms. One advantage that shouldn't be overlooked is *jEdit's* comprehensive help system – it's well designed, comprehensive, and very readable.

To make a long story short, *jEdit's* benefits are almost completely outweighed by its weaknesses. Its syntax highlighting is almost as poor as you'll find, it doesn't run very fast, it has a propensity for leaving menus floating in the air even when you change to another application, and its shortcuts are counter-intuitive.

VERDICT

Features	8/10
Ease of use	5/10
Documentation	9/10
Performance	6/10

A bit of a mess, with some power features thrown in for good luck.

LINUX FORMAT RATING
7/10

RoundupTextEditors

Emacs

Despite sounding like an Apple product, it's a remarkably powerful editor...

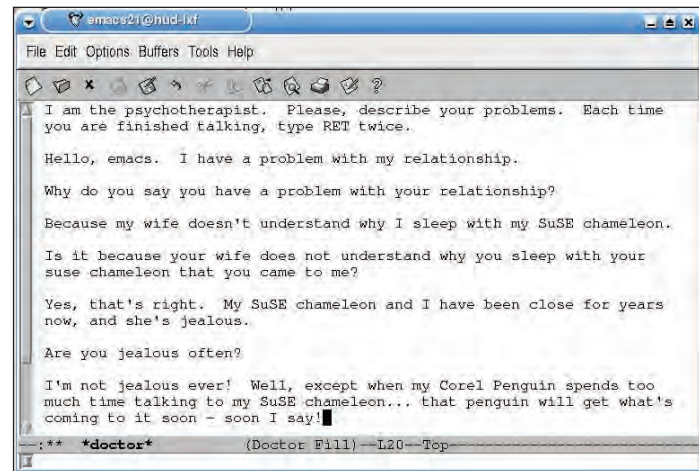
■ **VERSION** 21.3.2 ■ **WEB** www.gnu.org/software/emacs/

Emacs is renowned for being a combined text editor/kitchen sink package, although some people have yet to find the text editor in there. Originally created by Richard Stallman as a set of macros for the TECO text editor (EMACS stood for Editor MACroS), Guy Steele worked with RMS to turn it into a developer's text editor in its own right, and now *Emacs* has grown into a complete environment – you can edit text in there, sure, but you can also compile and debug your programs, apply patches, use CVS, read your mail, play chess, and even talk to a virtual psychiatrist (shown in screenshot).

Because Emacs has so many features, it uses a two-step command system – you press one command key followed by another to execute each action. For example, to quit *Emacs* you press Ctrl-X then Ctrl-C – not an

obvious combination to say the least, but once your mind switches wholly into the *Emacs* way of thinking, it starts to make sense. The situation is muddled somewhat because all too often *Emacs* using Ctrl + [key] to do something, and Alt + [key] to do the same thing in reverse – for example, Ctrl-V moves forward a page, whereas Alt-V moves back a page.

On the plus side, Emacs does everything you could want if you're a programmer as long as you're willing to spend the time learning it. Luckily, *Emacs* comes with a very comprehensive tutorial in many languages that guides you through learning the system from start to finish. Other popular advantages in *Emacs* are the scratch buffer, which is a handy place to keep text in between documents, spell checking, and its built-in *diff* tool to compare differences



between files. *Emacs* is capable of doing syntax highlighting, although it is admittedly clumsy to get working properly, and even then the final result is quite weak when held up against *Kate's* capabilities, for instance.

While its usefulness is undoubtedly restricted to advanced users, it's a great program, with lots of room for expansion and plugins that can increase its usefulness even more. As long as you don't mind spending a little time learning how to use it, then a lot of time actually remembering how to use it, then *Emacs* will certainly be of interest.

One half of the old "*Emacs* vs *Vim*" argument, it's favoured by power users, but loses out on keystrokes.

VERDICT

Features	9/10
Ease of use	7/10
Documentation	10/10
Performance	8/10


A better kitchen sink than most real kitchen sinks, with text editing to boot.

LINUX FORMAT RATING
8/10

Nano

As tiny as its name suggests, with a feature level to match – ideal for novices.

■ **VERSION** 1.2.1 ■ **WEB** www.nano-editor.org/



Nano takes 350 years of mathematical controversy in its stride!

Nano is GNU's replacement for the popular *Pico* text editor, and has much the same functionality. As editors go, they don't come much easier than *Nano* – all your commands are on screen all the time, meaning there's no doubt as to how to do anything. What *Nano* has in ease of use it is lost when it comes to functionality – it is a very basic text editor that allows you to, well, *edit text*, and not a whole lot more besides.

As such, it's perfect for all sorts of basic tasks, but becomes unusable for programming, complex editing, or anything beyond making simple changes to documents. When restricted to that use it's a fine editor – very fast, easy to get around, and easy to learn. While *Nano* does have some advanced features, such as syntax highlighting, they aren't anything like as advanced as what is provided by many other packages in the field.

Because *Nano* is a modeless editor – every character you press except control and escape are entered directly as text – it's much more natural to users of other operating

systems. When faced with *Vim*, *Emacs*, or even *Joe*, most new users get confused very quickly, which is where *Nano* has a firm lead. Not only is it easy to learn, but there are no levels of command – you can perform any command from any state in the program, meaning you can't "get lost" inside it.

The combination of very few commands, modeless editing, and "always-on" help mean that *Nano* is ideal for beginners looking to edit text and nothing else – quick to learn, easy to master, and generally available in every distribution.

VERDICT

Features	7/10
Ease of use	10/10
Documentation	7/10
Performance	8/10

Basic, yes, but is more than enough for most people just looking to edit text.

LINUX FORMAT RATING
8/10

TEXT EDITORS

THE VERDICT

With so many text editors available, it's unsurprising that the best six that we examine here are all very good. There's a strong mix of functionality in each of them – some focus more on the needs of programmers and advanced users (who are undoubtedly the primary user base for text editors) as opposed to word processors, but most offer some degree of functionality for everyone – easy file manipulation, project-level views, spell checking, and word wrapping.

At the bottom end of the scale are *Nano* and *Joe* – text editors, yes, and more advanced than many are commonly used in that they both support complex search and replace techniques. *Joe* has the edge on features by quite a margin, although it sacrifices usability a touch to do so. Both are very usable, and both are quick and easy to learn for everyone.

One step up, the debate between *Emacs* and *Vim* is likely to continue on for a long, long time – they are, after all, the only two serious options for “power user” terminal text editors. Neither *Nano* or *Joe* come close to *Vim* and *Emacs* when it comes to the amount of functionality offered, particularly for programmers, and as such the livelihood of *Vim* and *Emacs* is guaranteed for some time yet. *Emacs* is now looking more than a little out of date, and indeed out of *touch* with the way people work. The menu bar just keep getting longer and longer as more features are added without any substantial redesign, meaning that the learning curve is just getting steeper. *Vim*, on the other hand, continues to look deceptively easy-to-use, and is marching forward stronger than ever now that *KVim* and/or a *Vim KPart* is looking certain to be in KDE 3.2 – at last, we'll be able to use *Vim* to write our emails from *KMail*!

Leading the pack by no small margin is *Kate*, the primary editor of KDE. Apart from looking beautiful and integrating very well into the rest of KDE, *Kate* has features galore. Its left-hand pane for viewing open files and opening more files keeps your filesystem to hand at all times, and the ability to have a full *Konsole* shell open in the bottom pane is a dream come true. While it's not perfect – word counts on selection, please! – *Kate* still makes text editing a more enjoyable task. With a massive number of undos and redos, perfect syntax highlighting in more languages than most people have heard of, line numbering, power search and replace, and more, *Kate* is definitely worthy of our coveted *Top Stuff* award.

That just leaves *jEdit*, which, to be frank, almost seems to be floating in the middle of nowhere. Lacking much of the functionality of *Kate*, *Emacs*, or even *Vim*, *jEdit* is lagging behind the

game quite a bit. However, extensive work is underway to bring it up to speed – right now it does have one or two tricks up its sleeve that have yet to be seen in other editors (the excellent, if somewhat complicated, code folding is a good case in point), and it just needs to work harder on providing more basic core functionality.

Text editors are such a raw topic for many people, and I have little doubt that I'll be flamed something along the lines of “Ah, but you didn't do XYZ!” My only defence to that can be that if I, having used Linux for such a long time, had problems with certain things, then newcomers to Linux would certainly have no chance!

Regardless, all the editors have their own unique advantages, and there are dozens more we haven't covered due to space restrictions. Until next time, *Kate* is undoubtedly the shining light of text editors, and deservedly so. 

Ed: the standard text editor?

A piece of history

It would be very difficult to produce a comparison table for the half-dozen text editors we tested that contained much useful information. So instead we have a treat for fans of computing history! Following is a parody of the early “text editor war” messages sent around the Net. This classic is from the early 1990s, although sadly the author's name appears to have been lost in the mists of time...

When I log into my Xenix system with my 110 baud teletype, both vi “and” Emacs are just too damn slow. They print useless messages like, ‘C-h for help’ and “foo” File is read only’. So I use the editor that doesn't waste my VALUABLE time.

Ed, man!

!man ed

ED(1) UNIX Programmer's Manual
ED(1)

NAME
ed - text editor

SYNOPSIS
ed [-] [-x] [name]

DESCRIPTION
Ed is the standard text editor.

Computer Scientists love ed, not just because it comes first alphabetically, but because it's the standard. Everyone else loves ed because it's ED!

“Ed is the standard text editor.”

And ed doesn't waste space on my Timex Sinclair. Just look:

```
-rwxr-xr-x 1 root 24 Oct 29 1929 /bin/ed
-rwxr-xr-t 4 root 1310720 Jan 1 1970 /usr/ucb/vi
-rwxr-xr-x 1 root 5.89824e37 Oct 22 1990 /usr/bin/emacs
```

Of course, on the system “I” administrate, vi is symlinked to ed. Emacs has been replaced by a shell script which

- 1) Generates a syslog message at level LOG_EMERG;
- 2) reduces the user's disk quota by 100K; and
- 3) RUNS ED!!!!!!

“Ed is the standard text editor.”

Let's look at a typical novice's session with the mighty ed:

golem> ed

?

help

?

?

?

quit

?

exit

?

bye

?

hello?

?

eat flaming death

?

^C

?

^C

?

^D

?

Note the consistent user interface and error reportage. Ed is generous enough to flag errors, yet prudent enough not to overwhelm the novice with verbosity.

“Ed is the standard text editor”

Ed, the greatest WYGIWYG editor of all.

ED IS THE TRUE PATH TO NIRVANA!

ED HAS BEEN THE CHOICE OF EDUCATED AND IGNORANT ALIKE FOR CENTURIES! ED WILL NOT CORRUPT YOUR PRECIOUS BODILY FLUIDS!! ED IS THE STANDARD TEXT EDITOR! ED MAKES THE SUN SHINE AND THE BIRDS SING AND THE GRASS GREEN!!

When I use an editor, I don't want eight extra KILOBYTES of worthless help screens and cursor positioning code! I just want an Editor!! Not a “viitor”. Not a “emacsitor”. Those aren't even WORDS!! ED! ED! ED IS THE STANDARD!!

When IBM, in its ever-present omnipotence, needed to base their “edlin” on a UNIX standard, did they mimic vi? No. Emacs? Surely you jest. They chose the most karmic editor of all. The standard.

Ed is for those who can “remember” what they are working on. If you are an idiot, you should use Emacs. If you are an Emacs, you should not be vi. If you use ED, you are on THE PATH TO REDEMPTION. THE SO-CALLED “VISUAL” EDITORS HAVE BEEN PLACED HERE BY ED TO TEMPT THE FAITHLESS. DO NOT GIVE IN!!! THE MIGHTY ED HAS SPOKEN!!!

?

Hot Picks

The best new open source software on the planet!



Mike Saunders

A coder since Amiga times, Mike's a Linux and BSD guru.

This is the place where we get to profile some of the hottest software around.

Each month we trawl through the hundreds of open source projects which are released or updated, and select the newest, most inventive and best for your perusal. Most of the Hot Picks are available on our coverdiscs, but we've provided web links if you want to make sure you have the very latest version.

If you have any suggestions for things that we should cover, email us at linuxformat@futurenet.co.uk

Hot Picks at a glance

ReZound	38
Gestalter	39
Software Suspend	39
Madman	40
GNOME Commander	40
AIDE	41
Kahakai	41
Pathological	42
TuxPuck	42

HOT PICKS AWARD

Everything covered in our Hot Picks section is unmissable, but every month we'll be singling out one project for outstanding brilliance. Only the very best will be chosen!



SOUND EDITOR ReZound

■ VERSION 0.8.1 beta ■ WEB <http://rezound.sourceforge.net>



ReZound's main sample-editing window with the balance editor in front.

Linux already has an impressive market share in the server world, and is gaining more respectability on the desktop as each month passes. But while a great deal of effort is being put into friendly front-ends and easy-going apps, the professional application range is fairly limited at present. *The GIMP* (while not quite rivaling *Adobe Photoshop* in feature terms) was one of the first tools to demonstrate that Linux was far more than a hobbyists' toy; now, in the digital audio arena, we're starting to see a growing range of equally 'serious' software like *ReZound*, too.

Regular LXFers may be familiar with *ReZound* – issue 35 covered this audio file manipulation tool briefly in *Hot Picks*. Many months and releases later, much work has been done on the interface and heaps of new features have been added as it strolls towards a full 1.0 release. Also, with it gaining popularity we've had a few suggestions to give it another look and spread the word – as always, we value any ideas from you, our readers, on what to cover.

ReZound is applaudably simple to install; the developers have produced a statically linked binary package to eliminate dependency hassles (most of the required libraries are compiled right in). This makes for a larger package but as a result it's a walk in the park to get working. If you'd prefer to compile from source for tweaking or optimising, you'll need the *Fox* toolkit, *libaudiofile*, *Lame* and *Ogg* libraries installed – after that it's the usual build procedure.

File format support

Boasting an excellent range of supported sound file formats – WAV, AIFF, AU, raw, Ogg, FLAC and MP3 (via *Lame*) along with its own native .rez – *ReZound* can handle most popular formats in current use, although WMV is notably (and understandably) absent. By default it sends output to the standard OSS sound device, but it can also handle PortAudio and JACK, the latter being an ultra low-latency sound server and a pro's alternative to ESD and friends.

ReZound's competently designed and well-tooltipped *Fox* interface crams plenty in without falling into clutter territory. Beneath the control panel and open file list lies the sample pane which displays the channels. Particularly smart are the draggable wheels for zooming the display – much more elegant than a standard +/- widget, and is also supports the mouse wheel for scrolling through a sample.

Cutting, pasting and mixing between open samples is easy, and a multiple-undo option (including list of recent actions) is also present. Recording is a breeze and the time-limit setting is especially useful for unattended grabbing of lengthy radio content etc. On the whole, it's an approachable and pleasant to use UI with an adequate array of keybindings, although the supplied colour scheme may be a tad garish for some users.

ReZound allows for all manner of effects and filters to be used on samples, including volume change, reverse, gain, flange, distortion, 'IIR single pole' filters and 'biquad resonant' filters, and stacks more. Loop functions are well matched with the ability to add cues and make a selection symmetrical, while the usual resampling, normalise and graph-driven balance editor are all in the bundle. Lots of development work is going on and new code being put in place, so a CVS pull is the best way to grab all the latest features.

ReZound faces some competition in the form of *GLAME* and *Audacity*; the latter's slick interface and thorough featureset make it a better first choice if you're unfamiliar with audio editing tools, and *ReZound's* documentation is limited to a few decent but small online articles.

Still, it's fast, very solid (didn't flinch at our weighty 280MB files) and flexible, and with a bit more spit-shine it'll be another string to the Linux professional workstation's bow. Naturally it isn't completely as feature-packed as the specialist high-end tools, but if you do any kind of audio manipulation under Linux, give *ReZound* (and indeed *Audacity*) a look.

VECTOR GRAPHICS

Gestalter

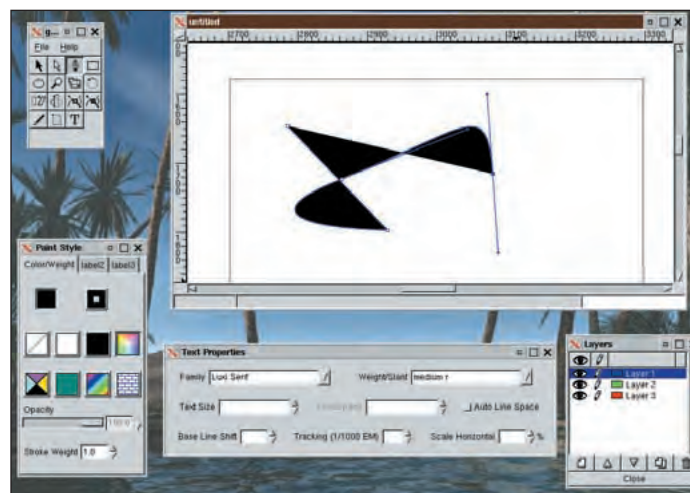
■ **VERSION** 0.6.4 ■ **WEB** www.linotux.ch/gestalter/

Long gone are the days when tools like *XPaint* were the high point of free graphics software – from its early releases *The GIMP* took centre stage, attracting masses of attention, and even Corel released its Wine-supported half-port of *PhotoPaint* on the OS. However, these have all focused on bitmap image editing, while vector-based packages have not been so quick to arrive. (Bitmap editors work with individual pixels, thus scaling-up poorly, while vector tools use co-ordinates, *a la* most clip-art.) *Gestalter* is a new project aiming to create a robust vector graphics program for Linux.

Because it's written in C++ and uses GNOME, you'll need the corresponding GTK+ and GNOME libraries (*gtkmm* and *gnomemm*) to

compile, along with *libstdc++*. These are available in most distros, although the pre-built RPMs are unusually picky about having the right versions – if you're confronted by 'unresolved symbol' errors, try downgrading to a slightly older *gtkmm* release.

As *Gestalter*'s front end shares similarities with Adobe's successful *Illustrator* program, it's not too daunting for newcomers and manages to save screen space with the packed-in main tool window. No keybindings for these functions appear to be in place yet, nor are any tooltips, but as the version number shows, it's early days. In addition to the essential Bezier curve function and basic shape creation tools, rotating, stretching, zooming and cropping are catered for, as is TrueType font text.



A bunch of the dialogs already implemented, along with the tool box.

With a layer implementation, Postscript file or direct printer output and rendered/wireframe viewing mode, *Gestalter* has a lot already in place although many features have yet to see any work and some menu items simply don't respond. The program saves in SVG format; sadly though, at present it can't open files in other formats, and it suffered when fed many other SVG files.

Gestalter is an intriguing project – young at heart with its hamster-photo About screen, yet it's just about usable for basic tasks and with a few more developers on board, it could turn out to be something very special indeed. With proper documentation, stability effort and general exposure, *Gestalter* is certainly something to keep a close eye on for future developments.

SOFTWARE HIBERNATION

Software Suspend

■ **VERSION** 1.0-pre15 ■ **WEB** <http://sourceforge.net/projects/swsusp/>

Recent laptops come with a 'hibernate' feature built into the machine's power-management functions. As opposed to the traditional power-saving feature of spinning down hard drives and keeping only the RAM going, hibernation writes the entire memory to disk and halts the machine – perfect for changing batteries and avoiding long restarts. Once powered up again, the system reads from disk and is (hopefully) in the exact same state. Linux's APM implementation can do this on supported hardware, but *Software Suspend* (*Swsusp*) allows it on almost any machine, laptop or not.

To get *Swsusp* working, you'll need to apply a small patch to your kernel; recompiling the kernel

requires planning, but we've had guides in *LXF* (issue 39, page 54) and

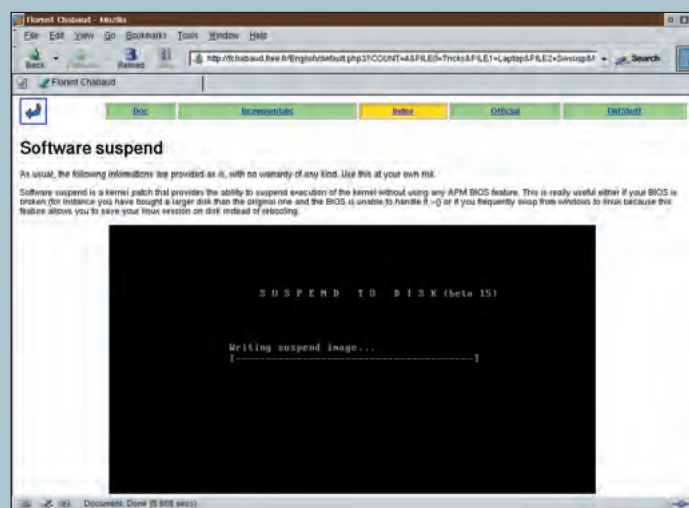
there's plenty of info on the Web. If you're unsure, it's best to wait until it is part of the official source tree. Alongside the patch is a script which ensures that devices and services are halted and re-initiated correctly – just run it directly to install it.

After the kernel patch has been applied, you'll need to enable `CONFIG_SOFTWARE_SUSPEND` in your `.config`, rebuild and install. *Swsusp* uses a swap partition to store the memory image, so an extra parameter

needs to be added to the kernel line of your bootloader, indicating that partition: `resume=/dev/hda3` or similar. Note that the current release only works on IDE drives for suspending, with SCSI support due in the near future.

After rebooting with the new kernel, parameter and script in place, calling `hibernate` as root begins the process – some picky modules are removed and services stopped, the text-mode progress-bar shows the memory being written to the swap partition, and then the machine powers off. On next boot, the kernel spots the swap partition's special signature and immediately starts filling up the memory.

We've been running *Swsusp* on a ThinkPad laptop and Dell desktop – the former works perfectly, as laptop hardware is usually more tolerant of quick stops and starts, but with some tweaking the desktop box is happy too. *Swsusp* is well worth investigating if your laptop doesn't support hibernation, you sleep in the same room as a noisy PC or you're simply tired of long shutdowns and reboots, and then restarting all your applications. Highly useful.

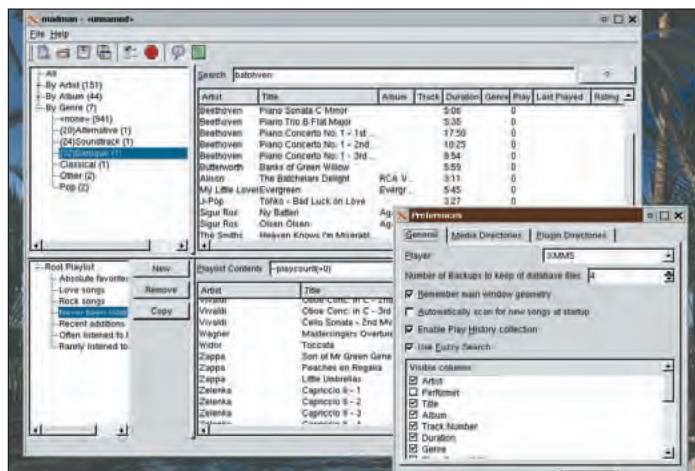


The screenshot on this site is, visually, about as exciting as it gets...

MP3/OGG JUKEBOX

Madman

■ VERSION 0.9.1 ■ WEB <http://madman.sourceforge.net>



Can't remember how to spell the third-best German composer's name? Don't remember where you've stored some music? *Madman* helps out.

Jukebox-esque music players are all the rage at the moment, helped in part by the influence of

Apple's *iTunes* software and the overall growth of hard drive sizes and online sharing communities. A couple of

Ogged or MP3ed albums may be manageable; extend this to an entire collection of hundreds or thousands of tracks, and it can become a colossal nightmare to sort out. We've recently looked at *Rhythmbox* and *JuK* in *Hot Picks* over the last few months, and now *Madman* joins the battle with a neat interface and impressive assortment of features.

To compile *Madman* from source, you'll need at least *Qt 3.1* (note that it's not a KDE app, though), along with the Ogg Vorbis libraries, *id3lib* and *XMMS* for playing. These should be already installed with all modern distros, making the build process much more problem-free than in the past.

When first started, *Madman* prompts for a directory of MP3 or Ogg files – it then reads the tags and presents a list as part of the four-paned display. The upper panes display the entire collection and an album tree for narrowing down, while the lower panes assist in selecting by popularity and genre. Simple and fresh, then, but it feels very cluttered at lower resolutions.

Tracks can be sorted by name, duration and time of last play, while ratings can be applied through the context menu. Fortunately, the displayed info columns are quite easily configured to your requirements. *Madman* currently uses *XMMS* to play selected songs – it's not particularly elegant to have a whole separate application pop up each time, but the developers are considering a built-in player in future releases (and apparently a back-end for others can be added with some quick coding).

Madman permits tag-editing on the fly (using a Windows-like slow double-click on the track), while the 'fuzzy search' feature is a superb addition for locating files when you're not sure of the exact name. Equally, the supplied CD-writing plugins, while not perfectly integrated, are nice touches too. Overall, it's a zippy and straightforward little jukebox, and deserving of a try if you just want the essentials sans superficial frills.

If you'd prefer to build your own MP3 jukebox program, read our tutorial on page 60 to find out how.

FILE MANAGER

GNOME Commander

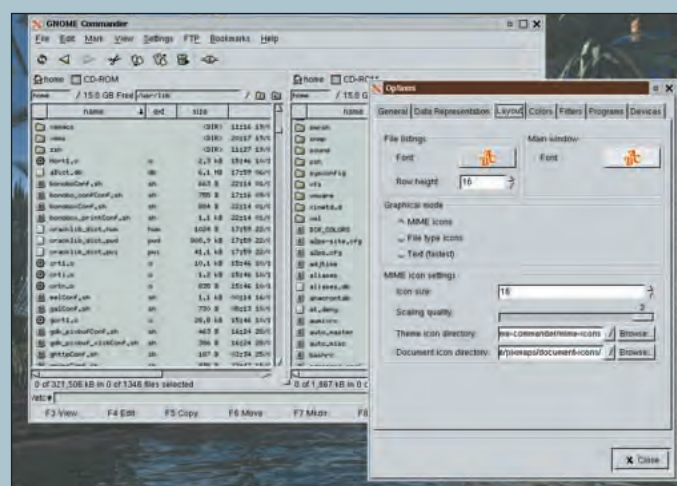
■ VERSION 1.0.1 ■ WEB www.nongnu.org/gcmd/

File managers, like desktops and text editors, typically provoke aggressive debate and flame-wars – everyone has their own personal preference, and opinions on what is right/wrong never cease. Some want flashy built-in previewers and HTML rendering engines, some can't bear anything other than a thin wrapper around the *Bash* prompt, and others look for a middle-ground where imagery and effects don't kill speed and stability. *GNOME Commander* (*GCMD*) strives for performance and power-user affinity, while still being reasonably accessible to newcomers.

With the current releases built around *GNOME 1.4*, you'll need the relevant versions of *gnome-libs*, *gnome-vfs* and *gdk-pixbuf* to compile

correctly. Most distros which include *GNOME 2* still have the old packages available, and the team of coders behind *GCMD* hopes that a version for the new suite will be 'somewhere down the line'. We hope so too, as *GCMD* could one day prove to be an appropriate replacement for the snazzy-but-heavy *Nautilus* on lower-specced boxes.

GCMD's interface is modeled on the popular *Norton Commander* layout (as also used in the text-based *Midnight Commander*), containing a two-paned display of the filesystem. Along with the small toolbar for rudimentary navigation operations, a single-line shell prompt is provided for those times when clicking around is too inefficient – simple to use, then,



The default (somewhat headachey) dark-blue background for the main window has been replaced with a saner grey here.

although the traditional white-on-blue might not be that kind on our eyes.

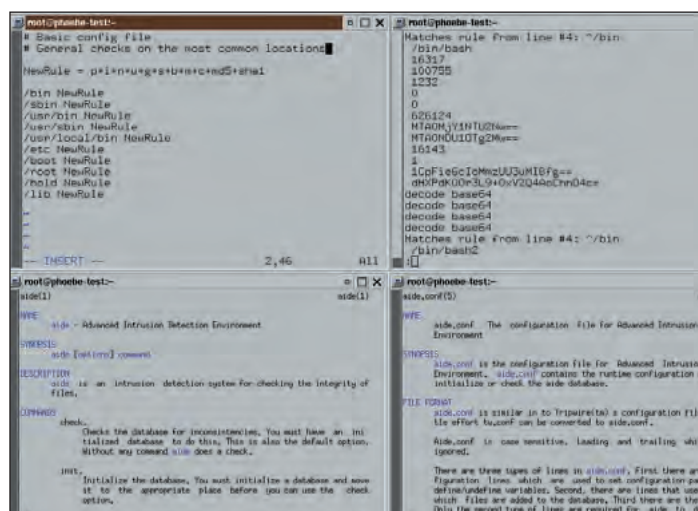
The usual copy, move, rename, delete and symlink functions all work well, as does the drag-'n'-drop. There's a detailed batch-rename feature (saves writing a shell script) and permissions-editing box, while selections can be made using normal *Bash* syntax or full *regexps*. Pretty solid on the feature front, then, and bookmarks and FTP browsing are other nice additions.

GCMD isn't in-your-face and world-explodingly amazing; as a small, svelte and reliable file manager it delivers well and there are no annoying distractions. The thin documentation is acceptable and there are a good deal of configuration options, but all things considered it's best as a halfway point between eye-candy and the raw shell prompt – just what many users with evolving abilities need.

SERVER INTRUSION DETECTOR

AIDE

■ VERSION 0.9 ■ WEB www.cs.tut.fi/~rammer/aide.html



Config file, verbose output showing its workings as it compares files, and the man pages.

Perhaps the most important job of any server administrator is maintaining security. There's no use in having a rock-solid uptime-

happy box crammed with the latest and greatest tools to wring maximum efficiency from your setup and give your users massive productivity

facilities if some script kiddie on the Net can break into it at any point. Even with a strong security policy in place and regular patches applied, breaches can and do occur, and this is when utilities like *AIDE* come into play.

Standing for *Advanced Intrusion Detection Environment*, *AIDE* helps the administrator after a break-in by listing potentially tampered files.

Flex and *Bison* are required for compilation; these are installed on most boxes with the development toolchain; if you don't have these, *LXF* generally includes them on our coverdiscs in connection with the ongoing *Compiler Writing* tutorial in *Linux Pro*. Additionally, the *mhash* library needs to be present as it provides a handful of extra checking algorithms (and it's also provided on our coverdisc). You'll need to build the static version of the lib, as (for safety reasons) *AIDE* compiles into a non-dynamically linked binary.

Once built, a configuration file has to be created. *AIDE*'s main task is to keep track of specified directories – it performs numerous 'digest' calculations or checksums on the files contained therein to see if they've

been altered. These include MD5, SHA1, Tiger and others, storing individual 'fingerprints' of the files in a database, and ensuring that if just a single byte of a file changes, at least one of the algorithms will spot it.

AIDE's top-notch flexibility lies in the user-defined rules for checking directories: *regexps* can be used for detailed file matching, and permissions, size, ownership and m/a/ctime among others can be monitored. Locations like */bin* and */sbin* are naturally worth watching the most, as they'd be the first port of call for a cracker, and the author rightly suggests keeping the binary and its database on a read-only filesystem.

As a quick and relatively uncomplicated alternative to the likes of *Tripwire*, *AIDE* behaves well for keeping tabs on a system. Throw the report into a nightly email *cron* job, and you can feel happier that any trojaned binaries on a box will be spotted fairly quickly, or feel safer when doing a post-mortem after something untoward has happened. It won't make a system 100% bullet-proof – no program can do that – but it's a good step to take.

WINDOW MANAGER

Kahakai

■ VERSION 0.21 ■ WEB <http://kahakai.sourceforge.net>

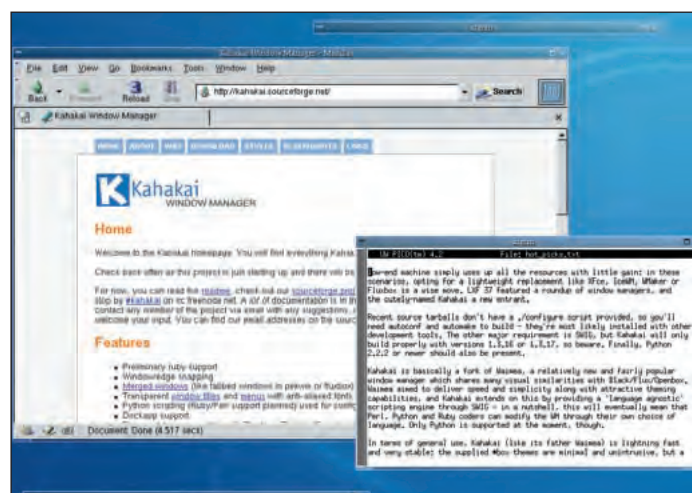
Even though KDE and GNOME continue to attract new users and eat up a good slice of the Linux desktop pie, there's still plenty of room for the more traditional window managers. Having a full desktop environment running on a low-end machine simply uses up all the resources with little gain; in these scenarios, opting for a lightweight replacement like *XFce*, *IceWM*, *WMaker* or *Fluxbox* is a wise move. *LXF* 37 featured a roundup of window managers, and the cutely-named *Kahakai* a new entrant.

Recent source tarballs don't have a *.configure* script provided, so you'll need *autoconf* and *automake* to build – they're most likely installed with other development tools. The other major requirement is *SWIG* (*Simplified Wrapper and Interface Generator*), but *Kahakai*

will only build properly with versions 1.316 or 1.317, so beware. Finally, Python 2.2.2 or newer should also be present.

Kahakai is basically a fork of *Waimea*, a new-ish and fairly popular window manager which shares many visual similarities with *Blackbox*, *Fluxbox* or *Openbox*. *Waimea* aimed to deliver speed and simplicity along with attractive theming capabilities, and *Kahakai* extends on this by providing a 'language agnostic' scripting engine through *SWIG* – in a nutshell, this will eventually mean that Perl, Python and Ruby coders can modify the WM through their own choice of language. Only Python is supported at the moment, though.

In terms of general use, *Kahakai* (like its forerunner *Waimea*) is lightning fast and very stable; the supplied 'box themes are minimal and unintrusive, but a bunch of sweet pixmap styles,



Kahakai in action, with its website open in Mozilla, and a clean theme.

including a smart Aqua-inspired one, can be added in. *Kahakai* will work with KDE and GNOME, has the usual virtual desktops and anti-aliased fonts, and the menus and keybindings can be altered by dipping into the text config files.

A fascinating new project, *Kahakai* will appeal to programmers – perhaps mostly to users of *Sawfish*, *SCWM* and the like – as a WM which provides a great deal of control via *SWIG* and its

multiple language feature. Once Perl and Ruby support have been added, *Kahakai* could well become one of the most popular lightweight WMs around.

Like "Wiki" (see the *LXF* website for more), "Kahakai" is a Hawaiian word. "Kaha" means primarily a scratch or mark, and "Kai" means the sea, or salt water. So "Kahakai" means the mark of the sea, or the junction or edge of the sea and land.

PUZZLE GAME

Pathological

■ **VERSION** 1.1.1 ■ **WEB** <http://pathological.sourceforge.net>

Periodically, a game comes along in which the concept appears so simple at first glance, but soon proves to be hellishly challenging. *Lemmings* is probably the best example of the 'Hah, this can't be

difficult' mindset rapidly deteriorating into 'Now that's just sick' through its crafty level design and general brilliance. Rainbow Arts' *Logical* was another such game and – as we've seen quite commonly in *Hot Picks* – Linux coders are doing a praiseworthy job of cloning and re-implementing the best game ideas and bringing more variety to the platform. In this case, *Logical* has become *Pathological*.


Both a recent version of Python and its corresponding Pygame library is needed to run *Pathological* – it's one of the new breed of interpreted language games, as we've seen before, and makes for faster development and more stable code from the off. A Windows version can also be found.

Pathological is well presented; it's marvellously polished and indeed looks better than a lot of shareware or commercial puzzlers. There's only one play mode, but the usual options to disable sound effects or the *Diggers-meets-Doom* in-game music are present, as is a fullscreen mode and high score board.

How does *Pathological* work? Not familiar with *Logical*? If you have ever played the 1970s MB Games classic

Downfall as a child, you will be immediately familiar with the concept behind *Pathological* – at its most basic, it's about moving coloured counters/balls around by turning a series of wheels. Essentially, you're in control of a bizarre electronic plumbing system, where coloured balls must be dropped in rotatable wheels. Get four of the same colour together, and that wheel is done. The problems come from the ball order, time limit and switches along the way – they can change a ball's colour, reject it, transport it, eat it and cause all sorts of other headaches.

Even with the hypnotic clicking of the balls and sometimes tough design of the 50 supplied levels, ultimately nobody else but the player feels responsible for any mistakes and that's the mark of a true classic. Clean, unambiguous graphics coupled with appropriate sound samples and quiet background tunes make it a first-rate sphere-'em-up. It won't appeal to everyone, admittedly, but it's certainly prepared to eat up hours of your time. Fear it.



Wowzers. And this is fairly straightforward compared to later levels...

AIR HOCKEY SIM

TuxPuck

■ **VERSION** 0.8.2 ■ **WEB** <http://home.no.net/munsuun/tuxpuck/>

A small seaside town past its heyday, a dingy arcade on the vacant sea front, an old beeping Tuppenny-Falls machine and a retirement-begging phosphor-burned Pole Position cabinet. A grumpy unapproachable man sits smoking in the Change booth, reading the Daily Sport. The floor is tacky. And in the corner is a heavily abused Air Hockey machine, waiting for another 50p coin and some bored kids to settle a score or shatter something with the surprisingly hard puck.

This is the modern image of Air Hockey (not helped by the lame flipper-equipped table-top toy versions which never, ever produce enough air to adequately float the puck), but it still has hardcore fans and is even considered a professional sport. 'Pucks in motion have been clocked at 81

mph and above!" brags one Air Hockey fansite. If you need to hone your skills on the table then *TuxPuck*, an Air Hockey sim featuring everyone's favourite loveable penguin mascot, could come to the rescue.

As with many of the games we've looked at in recent *Hot Picks*, *TuxPuck*'s developers haven't gone overboard on complexity – it's cinch to install, with no obscure dependencies. SDL, OggVorbis and FreeType2 are all included with recent distros, so compiling should be error-free.

TuxPuck opens with a polished intro (can be skipped), although you need to click the mouse to bring up the main menu – not entirely obvious, as the mouse pointer is never displayed or used in the menu. Currently there's only one option – play a match – but two opponents are provided: a sinister



Lovely Arcana's reflexes are no doubt hindered by the bottle of Gin (left).

blue Tux and a mystic maiden called Arcana. Tux is extremely fast and the tougher of the two, but Arcana's bizarre magic tricks on the table help her game out.

In play, the dark but well-drawn background and sprites are perfectly complemented by crystal-clear sound effects – the soft tapping of the puck

and harsh smashing of the glass work well, although a soundtrack sorely needs to be added. *TuxPuck* is clearly limited by the shallow game mechanics of Air Hockey; still, using the mouse feels 'natural' and requires speedy reflexes. Fun for a quick knockabout now and then, and you're unlikely to break your fingers... **LXF**

Better SOUND

Is Linux making the most of your soundcard? Whether you want improved game sound effects, or require high-quality music playback and MIDI recording capabilities, **Richard Drummond** investigates how to improve your Linux box's audio capabilities.

Sound is crucial in the role of the PC as a platform for delivering multimedia content. Whether you want to play a game, listen to Internet radio, watch a DVD movie or engage in video conferencing, your PC requires the ability to playback (and record) audio streams with high quality. If Linux is to be suitable for any of these applications and so move from the server farm or hosting facility to the desktop and the home, then it needs cutting-edge audio facilities too. Thankfully, these are beginning to appear, and, with a little work, Linux can be turned into a professional audio platform.

Most Linux distros today will routinely detect and configure a soundcard installed in your PC. But is your distro making the best use of your audio hardware? The chances are, if your requirements are modest, your existing sound set-up will suffice. But, if you want the best sound quality your system can deliver, you may need to delve a little deeper and tune your set-up yourself. If your interests are composing, sequencing or mixing your own music, then you'll definitely want to consider overhauling your distro's audio architecture.

So, how do you go about it? Well, that's what we'll look at in this article.


We're not going to concern ourselves too much with specific applications. You'll have to wait for a later issue for that, or refer to the Audio Editors roundup in *LXF38*. Instead, we're going to focus on the pieces of software infrastructure that you need to make your system sound its best.

Driving Linux audio

The Linux kernel officially gained support for sound hardware when Hannu Savolainen's soundcard driver was merged just prior to the release of Linux 1.0 in 1994. Savolainen's work quickly grew from providing basic support for Sound Blaster and similar



Sound



cards into a cross-platform audio framework for Unix-like operating systems. It was soon renamed the *Open Sound System* (or OSS – not to be confused with Open Source Software of course) and developed into a commercial product marketed by the company Savolainen co-founded, 4Front Technologies. He carried on maintaining the sound drivers in the kernel (which became known as OSS/Free) separately until Alan Cox took over in 1999.

While OSS/Free has been an adequate audio solution for Linux, it lacks the features required to turn Linux into a professional audio platform. Support for the latest soundcards has often been slow to appear, and in many ways OSS/Free has played second fiddle to its commercial sibling. Frustration at the shortcomings of OSS/Free led to the founding of the ALSA (Advanced Linux Sound Architecture) project in 1998 to create a brand new, open-source sound implementation for Linux.

The aim of the ALSA project is to replace OSS with a cleaner, more modular sound framework with drivers

to support a wide variety of audio adapters from the lowliest Sound Blaster to high-level, professional devices. ALSA provides a rich API, a flexible in-kernel sequencer for driving FM, wavetable and external MIDI devices, and a powerful plug-in layer which can manipulate audio streams on the fly. More than this, ALSA offers compatibility with OSS, through an in-kernel OSS emulation. While application support for ALSA is gaining, the ability to carry on using all the software that was written for OSS is an important point in ALSA's favour.

ALSA (0.9.4) is included on this month's coverdiscs. By the time you read this, the current ALSA release will be version 0.9.5 (see www.alsa-project.org/). This supports a much wider range of audio devices than the OSS/Free drivers, covering most of the popular ISA and PCI audio chipsets and cards available and even some USB and PCMCIA devices.

Multiple processor architectures are supported, including various proprietary and other audio devices on the ARM, PowerPC, PA-RISC and SPARC platforms. Older architectures

are less well represented, so if you're running Linux on an Atari ST or Amiga, for instance, you'll probably want to stick with OSS/Free. Also, support for some newer chipsets often isn't as complete as in the commercial OSS. For example, Creative's Audigy2 currently won't work with ALSA, but is supported by OSS (the SB Live!, Audigy and Audigy Platinum are all supported by ALSA, though). You can find out what cards and chipsets have ALSA drivers at www.alsa-project.org/alsa-doc/.

The commercial edition of OSS is still being developed and can be purchased from www.opensound.com/oss.html. An evaluation version can be downloaded for free. ALSA is definitely the future for audio on Linux, though. The ALSA audio drivers have been merged with the 2.5 kernel and will be the official sound solution in 2.6 when it is released. OSS/Free has been deprecated, and in fact is largely unmaintained in kernel 2.5. When the 2.6 kernel is released later this year, expect support for ALSA to accelerate – both from distro vendors and application developers. If you're

Anatomy of a soundcard

The evolution and capabilities of PC audio


A soundcard can use one of two fundamental methods to generate a sound waveform. Either the waveform can be created mathematically – synthesized – or a previously stored digital representation of the waveform can be played back via a digital-to-analog converter (DAC). The first method should be familiar to fans of pop music of the early 1980s, while the second is how sound is played back from a CD.

Early soundcards employed synthesis only. The card that launched the audio revolution on the PC, the 1987 AdLib Music Synthesizer Card, was based on a Yamaha OPL2 synthesizer chip which uses a technique called Frequency Modulation (FM). With FM synthesis, sound is produced by mixing (or modulating) pure sine waves to approximate more complex waveforms. The OPL2 uses two sine waves, known as 'operators', to generate each voice, and its total of 18 operators permits nine simultaneous voices. The results possible with only two operators are rather artificial but infinitely better than the beeps and buzzes of the internal PC speaker, so the AdLib was a huge success. Later FM cards improved sound quality by using the more advanced OPL3 synthesizer which supports four operators per voice and a maximum of 18 voices.

The AdLib became a *de facto* standard, but it is now largely forgotten thanks to Creative Labs' Sound Blaster (SB). The SB also had an OPL2 synth so was compatible with the AdLib, but, importantly, it added support for playing back and capturing digital audio, and thus could more accurately generate real instruments and sound. The method universally used for encoding waveforms digitally is called Pulse Code Modulation (PCM) and a digital audio unit such as the SB's is often called a PCM device. Some call it a DSP (Digital Signal Processor), but that term is often misleading. Some cards may have DSPs, which can perform mathematical tricks with digital waveforms to create various effects. The early Sound Blaster didn't, and its 'DSP' is just really a CODEC (coder/decoder) for converting PCM waveforms. In PCM encoding, the amplitude of a waveform is sampled at regular intervals – called the rate – and each sample stored as a number. The range of numbers that can be used to represent a sample is called the resolution. The higher the rate and the resolution, the more accurately the waveform can be represented. The first SB card had single channel (mono) sound with a resolution of 8 bits (each sample

encoded as a number between 0 and 255) and a maximum rate or frequency of 22050 Hz (22050 samples per second). In comparison CD audio employs stereo sound with each channel encoded with 16-bit resolution (a sample is between -32768 and 32767) at a frequency of 44100 Hz. Later Sound Blasters added stereo, 16-bit CD-quality sound and, eventually, more sound channels and even DSPs.

High quality digital audio requires a lot of storage space, so digital audio didn't replace synthesis overnight. The next technology to gain popularity was wavetable synthesis. While FM synthesis uses sine waves as the basis of sound, wavetable synthesis uses digital recordings of real instruments and sounds, which are then mathematically manipulated to generate the notes and effects required. In early cards, the instrument sounds were stored in ROM in the card, but later cards stored them in RAM – either on the card itself or in the computer's memory – so



new instruments can be uploaded. The sound quality of wavetable cards depends on the quality of the recordings

used, though. While the best cards can sound great, the cheaper ones are often no better than FM devices.

Current sound cards typically offer a mixture of these three technologies – either FM or wavetable synthesis and digital audio. Low-end devices, (eg those integrated with motherboard chipsets) often only support digital audio. Higher-end devices offer multi-channel digital audio, 3D sound, wavetable synthesis, DSPs, and digital audio outputs. Professional cards often provide higher resolutions, such as 24-bit or 32-bit digital audio.

not using ALSA yet, you'd better get yourself prepared.

Installing ALSA

ALSA consists of three main components: a set of soundcard drivers, which are provided as kernel modules; a shared library, *libasound2*, needed by client software that directly uses ALSA; and a set of tools. These three components are available separately from the ALSA website as source tarballs and are called *alsa-drivers*, *alsa-lib* and *alsa-utils* respectively. (There's a supplemental archive called *alsa-tools* which contains extra tools which may be useful to enable features of particular soundcards.)

If you're lucky, your distro already has ALSA support. Mandrake, SuSE and Debian do. Red Hat, notably, does not. However, you can get up-to-date pre-built ALSA binary packages for Red Hat 8.0 and 9.0 (and Yellow Dog 3.0) from <http://freshrpms.net>. There you can find ALSA kernel modules built for all these distros' default kernels. Get the *alsa-driver*, *alsa-lib* and *alsa-utils* RPMs plus the ALSA kernel modules package appropriate for your kernel. Once installed with your favourite RPM tool, you'll need to configure ALSA for your set-up. See the ALSA configuration section over on page 48.

If your distro doesn't support ALSA and you can't find pre-built packages elsewhere, then you'll have to compile ALSA yourself. Don't worry, it's not as difficult as it sounds.

The first task is to build the ALSA drivers you need. These are kernel modules which must be compiled against the kernel you wish to use them with, and you need a fully configured kernel source tree for this. It's not enough simply to use your distro's kernel source code package. The kernel tree must be correctly configured and **make dep** must have been run on the source tree. This is necessary for the ALSA modules to be able to correctly resolve dependencies against the kernel.

If you are running a default kernel from your distro, you need to find the config file that your distro vendor used to build that kernel. Often this will be installed in your /boot directory along with the kernel image itself. For example, if I'm using a kernel called 2.4.21-rich2 and I have the source

code installed in /usr/src/kernel-source-2.4.21 then I can configure this source tree ready for building ALSA as follows:

```
cd /usr/src/kernel-source-2.4.21
cp /boot/config-2.4.21-rich2 .config
make oldconfig
make dep
```

If you are building a custom kernel, then obviously you'll be configuring and doing a **make dep** during the process. When you configure your kernel, there are few things to be aware of. Firstly, ALSA requires sound support, preferably built as a loadable module. (You can also build any OSS driver you want as modules, but they cannot be used at the same time as ALSA).

You must also select any modules that are appropriate for building ALSA drivers for your hardware. Obviously, PCI drivers requires PCI support and ISA Plug-and-Play drivers require PnP support, but also remember to include USB if you have a USB device and PCMCIA support for PCMCIA soundcards. Depending on your architecture, some drivers may have additional requirements. For instance, the driver for the PowerMac built-in audio device needs I2C support. Also, if you want to use your soundcard's joystick port, you'll need build the gameport and joydev modules, as well as drivers for your specific controller. Some cards will also depend on the kernel driver for classic ISA gameports (the ns558 module).

Once your kernel source has been prepared (or compiled if it's a custom kernel), the actual building and installing of the ALSA drivers is straightforward. Simply unpack the *alsa-driver* package, enter its directory and run the configure script followed by **make** and **make install**. The most important configuration option is **--with-cards** that specifies which drivers to build. As usual, you can find out what configuration options are supported with **./configure --help**, and this will also show a list of available drivers. Either use **--with-cards=all** to build all the drivers suitable for your system or supply a comma-separated list of drivers. You can find out which drivers support which cards at www.alsa-project.org/alsa-doc. The **--with-kernel** option tells the build system where



the kernel tree that you are compiling against is located. The default is /usr/src/linux, but if yours is elsewhere supply its path. Other important options are **--with-oss** which enables the OSS emulation, which you'll definitely want, and **--with-sequencer** which enables MIDI support. If your card has an FM or wavetable synthesizer or a MIDI interface, you want this. In fact, even if you don't have MIDI hardware on board, enable it because you can then hook up a software MIDI synthesizer to ALSA (more on this later).

The final step is to run the script */snddevices* to create the necessary device files to communicate with ALSA. You should skip this step if your system uses the device filesystem (*devfs*), since the drivers will create their own device files. This script sets the permissions of the devices it

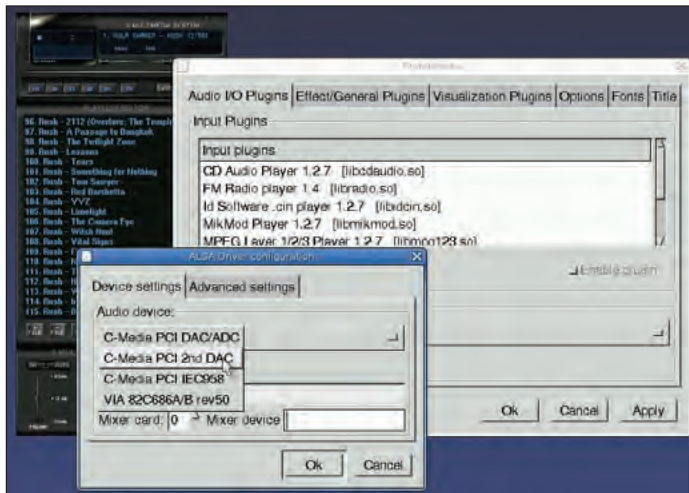
"If your distro doesn't support ALSA, you'll have to compile ALSA yourself. Don't worry, it's not as difficult as it sounds."

creates for root access only, so you'll have to modify these to let users access audio devices.

The following gives an example of how to configure, build and install a set of ALSA drivers for a CMI8738-based card (the **cmipci** driver), a Sound Blaster 64 (**ens1370**) and for VIA integrated chipsets (**via82xx**). Note you'll need to be root to do the install and modify device file permissions.



Sound



Application support for ALSA is increasing, and music players such as XMMS can output natively to ALSA.



```
tar xvf /root/alsa-driver-0.9.5.tar.bz2
--bzip2
cd alsa-driver-0.9.5
./configure --with-kernel=/usr/src/
kernel-source-2.4.21 --with-
oss=yes --with-sequencer=yes --
with-cards=cmipci,ens1370,via82xx
make && make install
./snddevices
chown root:audio /dev/snd/* /dev/sound/*
chmod 0660 /dev/snd/* /dev/sound/*
```

“It would be tedious to have to manually load ALSA drivers, so it’s a good idea to set up your system to do it automatically.”

OSS devices

The user space interface that OSS provides and ALSA emulates

The OSS API provides the following device files (all *char* devices with major number 14) for accessing the various functions on a sound card.

Device	Minor No.	Description
/dev/mixer	0	Mixer control
/dev/sequencer	1	Music playback (via internal synth or external MIDI device)
/dev/midi00	2	First MIDI port (sometimes called just /dev/midi)
/dev/dsp	3	Digital audio playback and capture
/dev/audio	4	Sun-compatible digital audio interface
/dev/sndstat	6	Sound card status information
/dev/sequencer2	8	Alternate sequencer interface (sometimes called /dev/music)
/dev/dmidi	9	Raw access to first MIDI port
/dev/dmfm	10	Raw access to internal synthesizer device
/dev/adsp	12	Second digital audio channel

A second soundcard will have similar device files with minor numbers starting at 16 with names /dev/mixer1, /dev/midi10, /dev/dsp1, etc.; a third card will start at 32; and so on. Exceptions to this rule are the devices /dev/sequencer, /dev/sequencer2 and /dev/sndstat which are global interfaces rather than per individual sound card. For example, the single /dev/sequencer interface provides access to all the synth or MIDI devices attached to your system.

Now all you need to do is build and install the *alsa-lib* and *alsa-utils* packages. This should just require a simple **./configure && make && make install** in each case. Use the **--prefix** configuration option to select the path under which the binaries will be installed.

Configuring ALSA

Now that ALSA is installed, it’s time to set things up. Actually, very little configuration is necessary. You can simply use **modprobe** to insert the ALSA modules you need. For instance,

```
modprobe snd-via82xx
```

will load the driver for my VIA chipset and any other modules that it requires, and this device will then be ready for use. The base module for each driver’s chipset is always named in this manner with the prefix **snd-**. Some drivers may require configuration options when loading. For instance, for old non-PnP ISA devices, you’ll have to specify the port, IRQ and DMA channel for the driver to use. See your card’s ALSA documentation for more information on what options the driver supports. You can also get a summary of a module’s options with the **modinfo** command.

It would be rather tedious to have to manually load the ALSA drivers and remember any options necessary each time you boot, so it’s a good idea to set up your system to do it automatically. The first step is to add some lines your /etc/modules.conf file to tell the kernel module loader about ALSA. For example:

```
# General ALSA options
alias char-major-116 snd
alias char-major-14 soundcore
options snd major=116

# Native config for card 0
alias snd-card-0 snd-cmipci
options snd-cmipci fm_port=0x388
mpu_port=0x330

# OSS config for card 0
alias sound-slot-0 snd-card-0
alias sound-service-0-0 snd-mixer-oss
alias sound-service-0-1 snd-seq-oss
alias sound-service-0-3 snd-pcm-oss
alias sound-service-0-8 snd-seq-oss
alias sound-service-0-12 snd-pcm-oss
```

The first section above supplies the global configuration, and you can

insert this verbatim without worrying too much about it. The next section identifies the first soundcard in the system and specifies any options. The **alias** statement says that when somebody asks for the first card **snd-card-0**, load the module **snd-cmipci** – which is the driver for my CMedia card. Simply substitute the name of your driver here.

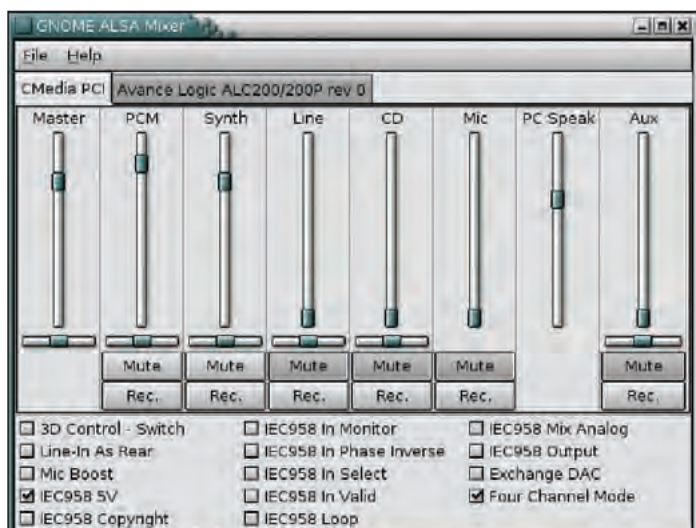
The **options** demonstrates how to supply any options that a driver requires. The next section sets up module-loading for the OSS emulation, and again you can just copy this as is this into your config. **sound-slot-0** is what OSS calls the first sound card, and the various **sound-service-0-x** entries set up mappings for the various OSS functions, where **x** represents the minor number of the OSS device being emulated. Thus **sound-service-0-0** represent what OSS would call the mixer device on the first soundcard. This alias will cause the **snd-mixer-oss** module to be loaded when the /dev/mixer device is accessed.

If you have additional soundcards, they need their own definitions in /etc/modules.conf, too. For example, my second card is described with the following excerpt. Note, this time, we only declare aliases for OSS mixer and PCM services, because you cannot have more than one /dev/sequencer device.

```
#Native config for card 1
alias snd-card-1 snd-via82xx
pre-install snd-card-1 modprobe
snd-card-0

# OSS config for card 1
alias sound-slot-1 snd-card-1
alias sound-service-1-0 snd-mixer-oss
alias sound-service-1-3 snd-pcm-oss
```

The above settings allow demand-loading of your drivers only when the OSS API is accessed – not the native ALSA API. Why bother, then? Well, you need to specify your driver options, and, anyway, these setting are needed by ALSA’s *init* script, which will have been installed in /etc/init.d/alsasound. You’ll need to use a run-level editor to ensure that this script is called appropriately during boot-up.



A mixer that supports ALSA natively, such as the **GNOME Alsa Mixer**, will give you more control over your sound card.

If you want to test ALSA now, **modprobe** your driver or execute `/etc/init.d/alsasound`. Examine the file `/proc/asound/devices` and make sure all the expected devices are listed. Now load the OSS emulation layer with **modprobe snd-mixer-oss snd-pcm-oss snd-seq-oss** and view the file `/proc/asound/oss/sndstat` (or `/dev/sndstat` – it's the same thing) and check that all your devices are recognised by the OSS emulation.

Digital audio

Before you can start making noise with your ALSA setup, you'll need to run a mixer tool to turn up the volume. When first loaded, all the channels in the mixer will be muted. Run **alsamixer** or any other mixer and crank up the dials.

Although an OSS mixer client will work fine, you'll probably want to use a mixer that supports the ALSA API directly, since ALSA may provide extra controls via the mixer besides just the volume, mute and capture controls provided by OSS. The **alsactl** command can be used to save your mixer

settings to a file (by default `/etc/asound.state`) by executing **alsactl store**. The standard *init* script will restore these settings on bootup.

If you have the OSS emulation loaded, you can use the device files `/dev/dsp`, `/dev/adsp`, etc to play back and record sound to and from your soundcard. ALSA does provide its own set of tools for the job, though – **aplay** and **arecord**. If you use these without any arguments, then they will access the default PCM device, the first digital audio device on the first soundcard, and they will, respectively, play or record audio streams to the console. They will both work with raw audio data, WAV files or Sun audio. For example:

```
aplay /usr/share/sounds/pop.wav
```

To specify a different PCM device to play to, use the **-D** switch followed by the device name (see box above right *ALSA devices*). Running **aplay -l** will list all the valid digital audio playback devices on your system. Thus

```
aplay -D pcm.hw:0,1 sound.au
```

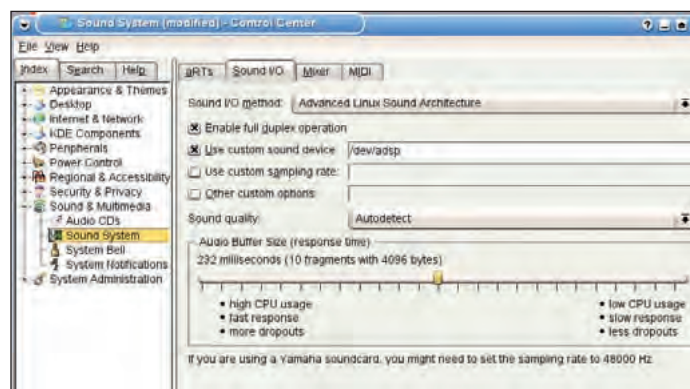
will play the waveform `sound.au` to the second device on the first card (equivalent to `/dev/adsp`). The tool **arecord** can be used similarly, and, again, **-l** will list valid capture devices, while **-D** will specify the device to capture from. Use the **-t** switch to select the file format of the captured waveform (WAV is the default), **-c** the number of channels, **-r** the rate and **-f** the sample encoding. The command

ALSA devices

A summary of how ALSA identifies your audio devices

ALSA device	OSS equivalent	ALSA name	Description
<code>/dev/snd/controlC0</code>	<code>/dev/mixer</code>	<code>ctl.hw:0</code>	Mixer on first soundcard
<code>/dev/snd/controlC1</code>	<code>/dev/mixer1</code>	<code>ctl.hw:1</code>	Mixer on 2nd soundcard
<code>/dev/snd/pcmC0D0p</code>	<code>/dev/dsp</code>	<code>pcm.hw:0,0</code>	First digital playback device on first card
<code>/dev/snd/pcmC0D0c</code>	<code>/dev/dsp</code>	<code>pcm.hw:0,0</code>	First digital capture device on first card
<code>/dev/snd/pcmC0D1p</code>	<code>/dev/adsp</code>	<code>pcm.hw:0,1</code>	2nd digital capture device on first card
<code>/dev/snd/pcmC1D0p</code>	<code>/dev/dsp1</code>	<code>pcm.hw:1,0</code>	First digital playback device on 2nd card
<code>/dev/snd/midiC0D0</code>	<code>/dev/midi00</code>	<code>rawmidi.hw:0,0</code>	Raw access to MIDI port on first card
<code>/dev/snd/midiC1D0</code>	<code>/dev/midi10</code>	<code>rawmidi.hw:1,0</code>	Raw access to MIDI port on 2nd card

You generally don't need to concern yourself with ALSA device files. The ALSA library uses names such as those to shown in the third column above to locate audio devices. This scheme in general is **<type>.<plugin>:<card number>, <device number>, <subdevice number>** where **<type>** is the type of the device, such as **CTL** (a mixer) or **PCM** (digital audio); **<plugin>** specifies the plug-in method which ALSA uses to access the device; and the **<card number>**, **<device number>** and **<subdevice number>** uniquely identify the device being accessed. The method **hw** gives direct access to the underlying device. ALSA's plug-in layer is very flexible and plug-ins let you manipulate the data being sent to a device in various ways. For example, it will let you convert the sample rate or encoding of waveform being sent to a **PCM** device on the fly. You can also create your own virtual devices, which are basically aliases for such device specifications. Using a virtual device, you can create a **PCM** device that will perform operations on the waveform being played (or captured) according to the plug-ins you have specified allow. However, a full discussion of the plug-in system is beyond the scope of this article. See the *alsa-lib* documentation for more details.



```
arecord -D pcm.hw:1,0 -r 22050 -t raw -f S16_BE mysound.raw
```

would record a waveform from the second soundcard and output it as a raw stream called `mysound.raw` using signed 16-bit big-endian encoding at 22050 Hz. The option **-f cd** is a shortcut for specifying CD quality sound. Read the contents of the *aplay* and *arecord* man pages for a full list of options and features that these commands support.

Sound servers such as **artsd** enable many client programs to make noise at once.



Sound

What is MIDI?

The Music Instrument Digital Interface

MIDI is a hardware and protocol specification, that's been around since the early 1980s, for connecting electronic musical instruments to a controller or sequencer. At the hardware level, a MIDI interface is a serial device similar but not compatible with the ubiquitous RS-232 serial interface (the COM port in PC terms). At the protocol level, MIDI devices communicate by sending messages to each other. A basic MIDI message is the "note on" event, which tells a listening device to play a note at a particular pitch and loudness. MIDI devices communicate over 16 channels and a MIDI instrument number will be assigned to each channel. Thus a device listening to a particular channel will thus know with which instrument (or patch) to play notes with when receiving "note on" events with that channel. A patch is the collection of device settings to make the synthesizer produce a specific sound, such as an acoustic piano.

The MIDI standard specifies that devices should be able to play 128 pitched instruments, each able to at least 24 notes. Early MIDI devices however didn't agree on what patch corresponded to which instrument number. On one device, instrument 14 might be a trumpet, while on another it might be a xylophone. Later, General MIDI (GM) was specified with a standard 'patch map' for compatibility.

PC soundcards often provide MIDI compatibility; almost all at last have an external MIDI interface. Obviously, the storage and user interface capabilities of a PC make it a great MIDI sequencer. Moreover, the FM and wavetable synthesizers on many sound cards can usually be used for MIDI play back, when appropriate set of GM patches are downloaded. With a software synthesizer, even cards without an on-board synth can play back tracks via a digital audio device by generating and mixing the MIDI voices in software.



client streams together in software and outputting the result as single stereo stream (the Emu10k does something similar in hardware). This allows you, for instance, to hear sound effects in a game and listen to music on an MP3 player at the same time, while still allowing audio feedback from your windowing system.

A number of various sound server implementations have been created over the years. Two popular examples are *Esound*, the Enlightened Sound Daemon, which is employed by the GNOME desktop and *artsd*, the server supplied with the *Analog Realtime Synthesizer*, the streaming audio platform used by the KDE desktop (see www.arts-project.org). They are both readily available with current distros, so we won't cover installation here. *Esound* is launched by the GNOME desktop and can be configured in the *GNOME Control Center*; KDE and *artsd* enjoy a similar relation. Both can access your hardware via OSS or ALSA.

Sound servers such as *Esound* and *artsd* act as a layer over your audio hardware, and supply an API for clients to generate sound and tools for users. For example, *artsplay* can be used to play back waveforms, while *artsrec* will capture them (compare these with ALSA's *aplay* and *arecord*). *Esound* provides similar tools. Moreover, more sophisticated, GUI-based client software is available for both systems. *aRts* in particular boasts a powerful array of client software,

including *aRts Builder*, a GUI builder for hooking up *aRts* components, and *Brahms*, a MIDI sequencer.

The problem with a sound server is that, for it to be effective, all software that wants to play sound has to use the server's API for outputting sound. Since there's many different sound servers in use, and some software will want to use ALSA or OSS devices directly, this can be a problem. For example, when you launch a program that tries to access a PCM device while that device is in use by *artsd* or *Esound*, that program will block, waiting for whatever's hogging the device, the server, to let it go. *artsd* can be configured to suspend itself after a specified time (configurable in the *KDE Control Center*) and relinquish its lock on the audio hardware, thus letting 'legacy' software output sound. However, this means the next *aRts* client run will be delayed while *artsd* is restarted.

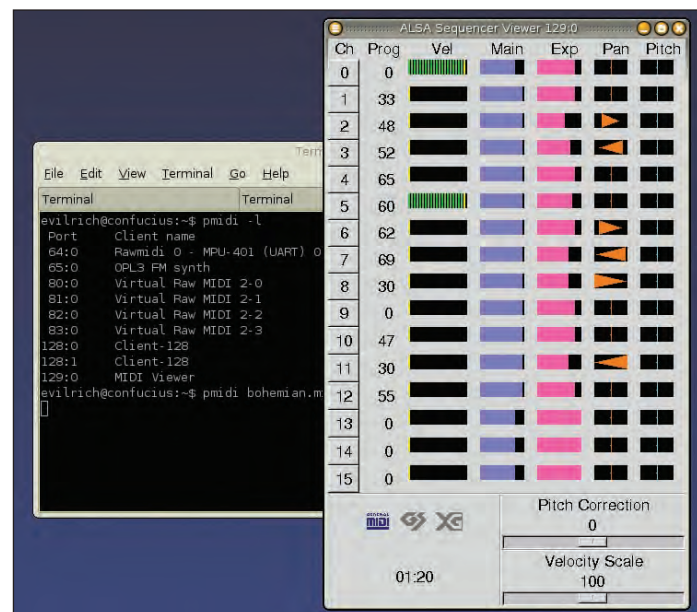
What happens if the server is in continual use, though? Well, typically, the offending program will block until you kill either it or the sound server. A rather drastic course of action! Both *artsd* and *Esound* provide an identical solution to this problem: a wrapper script – called *artsdsp* in the case of *artsd* and *esddsp* for *Esound* – that is used to launch the offending legacy program and redirect any access to */dev/dsp* to the server. It might sound like a kludge, but it actually works surprisingly well in practice. It's simple to use, too. For example:



Sound servers

Unless you have a DSP soundcard such as one based on the Emu10k chipset (the SB Live!, Audigy, etc.) then only one program can access a PCM channel at a time. This effectively means you can only play one stereo stream at a time. Many cards will have multiple PCM devices (perhaps to drive rear speakers for 3D sound and a digital output), so you may be able to play two or three streams, but on a multi-tasking desktop a transparent way of sharing digital audio resources between client programs is needed. The solution is a sound server. A sound server will let multiple clients output sound to a single device by mixing the various

RIGHT: ALSA's sequencer allows clients to create virtual MIDI ports to perform all kinds of neat tricks, such as this MIDI event viewer, *aseqview*.



esddsp /usr/games/ldoom

You can even use this solution to run multiple sound servers. Try this one and get the best of both worlds!

artsdsp esd&

An alternative, if you have multiple audio devices on your sound card, is to run the sound server on the second device (/dev/adsp in OSS terms) and so leave the primary device (/dev/dsp) free for other software. An even better solution is to have two soundcards and run the sound server on the second.

ALSA and MIDI

MIDI might have fallen out of favour with gamers, but it is still widely used by amateur musicians and in the recording industry, so it's no surprise that ALSA, as a professional audio platform, supports MIDI. ALSA drivers can make use of variety of MIDI-compatible hardware, including the FM or wavetable synthesizers on-board many common soundcards and external devices connected to the MIDI interface that most PC soundcards provide.

In ALSA, MIDI devices are identified by their port number. The ALSA sequencer works by routing MIDI messages between MIDI ports according to how these ports are connected and at the beat of its built-in timer. To get started with MIDI under ALSA, download and install the *pmidi*, an ALSA client for playing MIDI tracks (get it from www.parabola.demon.co.uk/alsa/pmidi.html). If you have an FM device, you also should get *playmidi* as well from (<http://playmidi.sf.net>). Although this is an OSS client, it does contain the GM patch files that you'll need to program your synth to play MIDI instruments – which *pmidi* doesn't. You'll also need the *sbiload* tool from

the *alsa-tools* package. If you have an AWE32/64 device, then you'll need the *sfxload* tool from the *awesfx* package to download soundfonts to your card. This appears not be supported by Creative any longer, but many distributions provide it as an RPM since it's needed by the kernel AWE driver.

When everything is installed, your first need to find the port number of your synth device. Run **pmidi -l** to list the MIDI output ports on your system. For playback on an internal device, you're looking for the port corresponding to your onboard synth. For example, my OPL3 device is labelled 'OP3 FM Port' and has port number 65:0. To use this for MIDI playback with *pmidi*, you now need to first download the MIDI patches. I can do this with

```
sbiload -p 65:0 -f /etc/playmidi/std.o3
/etc/playmidi/drums.o3
```

With an OPL2 synth leave out the **-f** switch and use the patches with the '.sb' suffix from *playmidi* instead. I can now play a midi track on this device with, for example,

```
pmidi -p 65:0 bohemian.mid
```

If you don't have a MIDI device, you can use a soft synthesizer. The latest version of the MIDI player *Timidity++* (on this month's LXF coverdiscs) can be run as an ALSA sequencer client to create a virtual MIDI port that will play back MIDI events via your PCM device. Simply launch *Timidity++* as a daemon with **timidity -iA -Os**

The **-iA** option tells *Timidity++* to run as an ALSA sequencer client, while **-Os** tells it to output to the default ALSA *pcm* device. Go to another terminal and run **pmidi -l** again. You should now see two ports labelled 'Timidity'. Find the port number of the



first, and use it with *pmidi* to play back MIDI tracks.

The ALSA tool *aconnect* can be used to connect MIDI ports to route messages between them. For example, say you have a piano-style

Freshmeat's *Timidity* project page has all the links you need.

“MIDI may have fallen out of favour with gamers, but is still widely used by musicians and the recording industry.”

musical keyboard plugged into MIDI interface. If you connect the input port of your MIDI interface to, say the output port of your FM device, you can use the keyboard to play notes on the FM device. List the input port on your system by running **aconnect -i**. If this is port 64:0, say, and the FM device is 65:0 as above, the you would execute

```
aconnect 64:0 65:0
```

to route message between the two devices. Use the same command with the **-d** switch to disconnect them when you are finished. **LXF**

Overcoming latency

How to improve the real-time performance of your Linux set-up

Playing and recording audio and MIDI sequencing is necessarily a real-time activity, not something that Linux has traditionally been good at. The problem is inherent latencies – due to the monolithic nature of the Linux and the design of Linux's task scheduler. A latency generally is the delay between an event occurring and the event being reacted to. An example is the delay between a task becoming ready to stream more data to

an audio device and the task being scheduled and thus actually sending the data. Such latencies can cause breaks and noise in audio playback and capture. A solution is to use large enough buffers to absorb any variations in scheduling behaviour. However, this causes its own latency – the time it takes data to move through a buffer.

Various solutions have been proposed for improving Linux's real-time response,

such as Andrew Morton's modification to improve scheduling latency (see www.zip.com.au/~akpm/linux/schedlat.html). This is available as a patch which must be applied to kernel source code and the kernel recompiled. Another such modification is Robert M Love's pre-emption patch, which lets a higher priority task pre-empt a task running in kernel routine – it effectively reduces the monolithic nature of the kernel. (see

www.tech9.net/rml/linux/)

Most end-users don't want to have to brave such tasks as patching the kernel, however. The good news is that the eagerly-awaited 2.6 kernel that will hopefully be with us by the end of the year contains many such modifications to improve real-time performance. If you can't wait until 2.6 is released, though, roll up your sleeves and read the Kernel Patching article in LXF29.

What on Earth is... RSYNC?

Nick Veitch gets fired up about delta compression file synchronisation.

>> What is this *rsync* then? It sounds like a tribute band, or, no, wait a minute, a special sort of pigment used when tattooing someone on...

No, I think I'll have to stop you there. Surprisingly enough, *rsync* is neither of these, but is a really useful little program for syncing files.

>> Syncing? So it's some sort of audio thing?

No. By syncing, I mean synchronising the content of files. Imagine you have two directories and you wanted the second one to always contain the same files as the first. *rsync* is a tool you could use to do this.

>> Er, right. That's nice. I can't actually think of any reason I'd want to do that though.

Well, there's more. Imagine the directories were on different computers, or different devices. You might have a set of documents, presentations or whatever on a desktop machine that you work on daily. Now you're going on a long trip and want to take them with you. *rsync* can be used to copy the files to your laptop.

>> Yes, I can sort of see that, but I could just copy the files couldn't I?

Well, yes, you could do that. But *rsync* is a whole lot smarter than that – it only downloads the files which are different. This can make a huge difference in the size of the transfer (obviously depending on the number and size of the files to be downloaded, and how often they change).

>> Hmmm, still not convinced it would be worthwhile...

Alright. Instead of copying files to your laptop, imagine you are downloading backup files from a server. This is actually a very common use for *rsync*. On our imaginary server, we keep five days-worth of backups.

We want to download these backups daily so we can keep a remote copy of them.

Using *rsync*, we can just grab the whole directory that the backups live in. As it checks the files, it will see that the previous four days-worth of files are redundant, and only grab the new one. This not only saves effort (it's much easier to set up), but also time and server load. These last two features are obviously pretty important.

>> Okay, but imagine I'm very foolish and don't keep backups, or I use some other system. Can I still go through my life in ignorance of *rsync*?

Let me give you another example of how *rsync* is useful. Last time you downloaded ISOs of your favourite distro, or even the Linux kernel source, if you were being clever and responsible, you probably did so from a local mirror site.

A mirror site has to contain exactly the same files as the original, otherwise it isn't worthwhile. But, how do the people running the mirrors know when the main site has changed? Do you think they check every day and look through the filelists to see what has changed?

The answer is 'No'. Mostly they use *rsync*. They could run a *cron* job to *rsync* against the parent site daily. If nothing has changed, then there is no damage done – the main server won't be overloaded by all the mirror sites trying to download the entire contents of the site every day. When files do change, only these are downloaded, reducing the amount of traffic. And best of all, the people running the mirror have no need to intervene personally and manually download. As long as the *cron* job is active, they need do nothing.

In fact, many sites operate an even more cunning system (the www.kernel.org site is one of them). When files are changed they send out a mail to registered mirrors. The mail can be sent to a special

account and used to automatically trigger *rsync*, so it only runs when the files change instead of on a daily basis. Cunning eh?

>> That does sound cunning. But how would you set up a script to read the mail? That sounds a bit tricky!

Well, surprisingly enough, you don't have to do too much work. If we stick with our example of www.kernel.org, you'll find an example Perl script in the `/pub/site/` directory. Other large sites have different systems for this sort of automation.

>> This is all very well, but what if my mirror isn't on a Linux server?

No problem. There is a Windows version of *rsync* available, so either the host or client or both can be Windows machines.

>> Okay, say I'm pretty interested in the idea now. Can you tell me who makes *rsync*?


rsync is a GPLed open source project written by, amongst others, the prolific Andrew Tridgell (who you may know from *Samba*, *wbench*, *Jitterbug* etc). However, there are dozens of active developers on the project, which is, as we have said, in widespread use.

>> Where do I get it then?

There is probably an *rsync* package available for your specific distribution. In many cases it is included in the base packages for your distro, so you may already have it. It is usually good to install from source if possible though – among other things it means that you know where you have installed all the important files!

The main site for the *rsync* source is ftp://rsync.samba.org/pub/rsync. There are some binary packages available in the binaries directory here if you want them.





WhatOnEarthRsync

« » I know you said previously that *rsync* is available on Windows. Is it available on other platforms too?

Yes it is. There are also different, compatible implementations of *rsync*, but also the whole protocol too. For example *jarsync* is a Java-based project to enable Java clients and servers to use the same delta-compression techniques and interact fully with *rsync* software. As it's Java-based, it should run on any platform supported by Java, which is quite a few.

This is a project very much still in development though. You can check out its progress at <http://jarsync.sourceforge.net>, but don't base your backup strategy around it just yet. There are various other Java-based projects too.

« » So assuming I'm going to get the source for Linux. Is it tricky to compile?

Not very difficult at all. Unpack the source archive and perform the standard:

```
./configure
make
make install
```

It isn't that big, so you should find it compiles in a few minutes. By default the binary will be installed in `/usr/local/bin`

« » Is this the server or client software?

It's actually both. *rsync* works as a server when run in daemon mode. For this to work properly, you also need to have an *rsyncd.conf* file, which we'll cover later. Now you have the client software, you need to know how to use it. The basic syntax is :

```
rsync [OPTIONS] [SOURCE] [DESTINATION]
```

There are loads of options, all of which are detailed in the man pages (type **man rsync** to read them), but for most jobs you will want to use the **a** (for archive, which actually sets a whole load of other

options), **z** which compresses the files for faster downloads, and **v** for verbose mode, so you know what is going on.

The source and destination can either be a local file or directory, or a network location (including *rsync* servers). In the case of the latter, this field could also contain a username (for logging in) and the possibility to specify a port number.

« » I'm sufficiently confused now. Hit me with some examples...

Fair enough. Let's take the UK mirror of the *Mozdev* site. *Mozdev* is the *Mozilla* developers' site, which contains all sorts of code, graphics and projects related to *Mozilla* and its various sub-projects. The first thing we can do is check what's available. Since this is an anonymous access server, we don't need to specify usernames or login. We can just do:

```
rsync -n mozdev.org.uk::
```

The **n** flag tells *rsync* not to transfer any data, just do a dry run. The consequence of this is that it lists all the files you would have downloaded. Running this example, we see only one entry, the *mozdev* directory itself. If we tack this on to the path we used in the previous command we can get a listing of that directory:

```
rsync -n mozdev.org.uk::mozdev
```

That's better, now we can see the dozens of directories that contain the files we want. Assume for a moment we want to create a mirror of the 'themes' directory. We'll use the switches discussed before to create a local archive of this directory, transferring the files compressed to save bandwidth:

```
rsync -zav mozdev.org.uk::mozdev/themes
/usr/local/mirror/
```

Note that this example assumes the directory `/usr/local/mirror` exists. The first time you run this command, *rsync* will download the complete contents of the directory, which may well take some time. Subsequently though, only the changed files will be retrieved. As *rsync* checksums the files, it will know which, if any have changed, and only download the relevant ones.

Note the **::** (double colon) after the URL here. This indicates that we are connecting to an *rsync* server, though you can also connect to a standard ftp protocol site, like the kernel site.

« » That's cool. How would this work on the kernel site then?

In pretty much the same way. For a full kernel mirror, you are supposed to fetch `/pub/linux` and `/pub/software`, which will consume a fair amount of disk space (and time, the first occasion you try it). If you want a mirror for personal use, you only need to get the bits you want. For example, to fetch the 2.6 kernel tree, you could use:

```
rsync -azv ftp.kernel.org::pub/linux/kernel/v2.6
/usr/mirror/kernel/
```

This would fetch just the **2.6** directory. When new versions are released, repeated running of the script



would fetch these too, but omit the existing versions, which obviously would not be changed.

« » Good stuff so far. Now what would I do to make sure it was always up-to-date?

Well, the easiest thing to do would be to run a *cron* script, and execute the same command every so often. For example, at ten past two in the morning or something (that's a good time by the way, you wouldn't believe the number of automated services that run exactly on the hour).

« » How do I create the config file for the server daemon then?

If you have ever meddled with Samba and its config files, you'll find the *rsyncd.conf* file pretty similar. The file is located by default as `/etc/rsyncd.conf` and contains some general settings first, followed by specific 'modules' for each filesystem path you wish to make available.

```
#/etc/rsyncd.conf
```

```
uid = nobody
```

```
gid = nobody
```

```
[backup]
```

```
comment = server backup files
```

```
path = /home/backup/data
```

```
[secret]
```

```
comment = Secret files, LXF only
```

```
path = /home/lxf/secret
```

```
hosts allow = *.linuxformat.co.uk
```


```
hosts deny = *
```

```
list = false
```

This config file illustrates a couple of useful things. The generic options at the start define the user and group ids of the *rsync* client when it connects. These are actually the defaults in this case.

All the global settings can be overwritten with settings in the individual modules. The module names are in square brackets. At the very least, they must have a path component. The comment is shown when the module is listed, as we did before.

As you can see from this config, the backup



module is pretty straightforward. Anyone selecting the module is given access (as user **nobody**) to the indicated directory.

In the second example, we have included a few more options to create a secret module. The comment is rather redundant as we are suppressing listing of this module (**list = false**).

We have also added a bit more security here, by only allowing clients from the **linuxformat.co.uk** domain to connect. This uses the fairly standard syntax for host settings – you can have multiple domains separated by commas, or use IP numbers if you wish. Remember to add the **deny = *** to deny other hosts!

» Is there any limit to the number of modules I can have?

No, not really, you can specify as many as you like, or have space for. It is probably worthwhile avoiding overlapping modules though.

» Now I have created the config file, how do I start the server?

To start *rsync* as a server daemon, you simply have to run:

```
rsync --daemon
```

Some versions of *rsync* don't come with an *init* script, so if you want to run this as a service at boot time, you might want to write your own. It's fairly simple. Here is an example that is taken from one of the *rsync* HowTos to give you an idea of what is required:

```
#!/bin/sh
# Rsyncd This shell script takes care of starting
and stopping the rsync daemon
# description: Rsync is an awesome replication
tool.

# Source function library.
. /etc/rc.d/init.d/functions

[ -f /usr/bin/rsync ] || exit 0

case "$1" in
start)
action "Starting rsyncd: " /usr/bin/rsync --daemon
;;
stop)
action "Stopping rsyncd: " killall rsync
;;
*)
echo "Usage: rsyncd {start|stop}"
exit 1
esac
exit 0
```

Copy this file to your usual *init.d* directory and make a symlink to the relevant *rc.d* *initlevel* directory. (for example, it may go in */etc/init.d/rsync* and be linked to */etc/rc3.d*).

» What about user accounts and security?

You can set up user accounts and use the *rsync* secrets file to control access. For real security though, it's best to access *rsync* through *ssh*, if you aren't creating a world viewable mirror. Using *ssh* makes it harder to take advantage of some of the security features of *ssh*, but does have the advantage of providing a known and secure access route without the need to open up any more ports or worry about more vulnerabilities.

To use *ssh* there is no need to change the *rsync* config, you merely have to specify *ssh* as the transport protocol on the command line. I think it might be time for an example again.

Assume we are using the server as set up with the config above. Imagine there is a user called backup on the system. We could download the 'secret' module with the following command:

```
rsync -avz -e ssh backup@linuxformat.co.uk
::secret /mirror/secrets
```

As you can see, this command is using the *avz* flags we used before, to specify an archive using compressed transfer.

The second switch specifies the protocol to use to contact the server, which we have specified as *ssh*. The destination address is given with the username prefixed to it with an *@* symbol, a fairly common convention. The double colon indicates that we are accessing the *rsync* server, and after that we put the module name. Finally, we add the destination path, in this case where we want to store the files on the local machine.

» When I try this I am prompted for a password. That's all well and good, but what if I want to automate the process?


You are prompted for a password because we are now using the *ssh* system. You could set up a user with no password, but in that case there's probably not a lot of point in using *ssh* in the first place.

It is still possible to automate this though, through the use of *ssh* keys. Generate a key on the client machine with *ssh-keygen*. The generated key can then be stored on the server in the users *.ssh/authorized_keys* file. Now when you log on, your local key is checked against the one of the server and the login proceeds automatically.

Now you can create a *cron* job on the client to automatically download the files at certain times in the same way as before.

» Woo-hoo! Okay, I'm impressed now. This does actually seem to be useful. Is it possible to run rsync on other platforms?

Yes. There are a number of ports available for various platforms. Perhaps the most interesting are the Java based ones, as obviously this gives a great degree of scope for running on many different systems,



including PDAs and embedded devices.

For Java, probably the most interesting project is *jarsync*. While not complete enough yet to be a thorough replacement for *rsync*, its development seems to be progressing nicely.

» What about MS Windows?

Yes, there are ports of *rsync* for Windows. In fact, most of these use the same code and the *Cygwin* environment. This will work on Windows 95/98/NT4/2000 and Windows XP. There are binaries available in various places, but probably the best location is http://optics.ph.unimelb.edu.au/help/rsync/rsync_pc1.html, which has links to binary downloads and instructions on setting up and installing on the Windows environment. The usage is pretty much the same, as on Windows you will run it in a shell (or you can set up batch files to perform repetitive tasks). On versions of Windows that support task scheduling, you could even set up regular backup retrieval tasks as we have already mentioned.


» Should I mention Mac OS X?

You can if you like. Actually, you'll find that *rsync* is included in the standard OS X release. You can also do all the other things we discussed such as using *SSH*, generating keys, running in daemon mode etc.

» And there's stuff for PDAs you say?

Yes. As we mentioned way back at the beginning when you were being so cynical, one cool thing you can do is effortlessly sync files on your desktop and your PDA, or for backing up data from such devices. For the Zaurus, there is a package called *zNetBackup2* (you'll find it on the Internet at www.dasgehtdichnichtsan.de/zaurus/smbmount.html) which will backup files to an *rsync* server on the network.

» Blimey. Thanks for pointing out another useful tool that I otherwise would never have noticed.

That's what this *What On Earth* section is for! Come back next issue to find out more about BSD... 

Tutorials >>

Our experts offer help and opinions on a whole host of Linux applications

YOUR GUIDE TO GETTING THINGS DONE!

Whether you are just starting out in Linux, or an experienced veteran, there's always more to learn. Every issue of *Linux Format* is packed full of practical advice, and nowhere is it more concentrated than in our tutorials pages.

Here you'll find expert guides to all sorts of things, from Basic Linux usage to understanding and deploying network solutions, from simple script coding to the complexities of Perl regular expressions, Java server apps and more. We aim to bring a good mix of tutorials to each issue, but if you have any suggestions for topics you'd like us to cover, why not contact us, by email at linuxformat@futurenet.co.uk or by snail mail, or log on to www.linuxformat.co.uk and post your suggestions in our special forums? Hope to hear from you soon!

Nick Veitch EDITOR

HOW CODE IS REPRESENTED

Including code in magazines can be tricky, but we hope our notation will help it become clear. When lines are too long for our columns, the remaining text appears on the next line in a solid blue box:

```
procedure
TfrmTextEditor.mniWordWrapClick
(Sender: TObject);
otherwise, there is usually a gap
between lines:
begin
mniWordWrap.Checked := false
end;
Often, you'll find the code on
the LXF CD/DVD too.
```

THIS MONTH TEACH YOURSELF...

MP3 jukebox

Making the MP3::Daemon Perl module earn its keep **p60**

Red Hat package management

Unofficial alternatives to Red Hat Network and Packaging Tool **p62**

Samba

Authentication and authorisation working in tandem **p70**

Realsoft 3D

Get started with the great 3D app on the coverdisc **p74**

Blender >>

First steps towards your own animated masterpiece **p78**

The GIMP

Integrating text: tools to use when a picture painting a thousand words isn't enough **p82**



ILLUSTRATION BY ROB DEBRICHY - www.blender3d.org

PHP

Optimisation of your code is no longer a black art – we show you some of the secrets **p86**

Beginners' KDE

Basic ins and outs of the Linux K Desktop Environment explained for novices **p90**

TIP OF THE MONTH! APT YOUR RPM!

Are you sick of reading about how great APT is compared to RPM? Have you yearned for the days when you no longer have to battle with dependencies just to install some software? Do you wish you could have the power of Debian's installation system without having to use Debian itself? Well, you can – visit www.rpmfind.net, and search for "apt" to get hold of the Debian APT package management suite for RPM-based distros, including Red Hat, Mandrake, and Yellow Dog. Once you've got APT installed, it works in conjunction with your existing RPM setup – it can search

for, download, and install the correct RPM files for your distro. There are two key commands to use with APT: **apt-get**, and **apt-cache**. The former is used to install and uninstall software, and the latter is used to query what's available for installation. To start using APT, try this command: **apt-cache search <package>** You'll need to replace **<package>** with the name of a package you want to search for. You can use regular expressions here, so, for example, **^lib** would search for all packages starting with "lib", such as *libwine*, *libbonobo*, etc. The chances are you'll get quite a long list back

from **apt-cache**, so either pipe it through **less** or through **grep** to get a filtered list.

To install packages once you've found them, use this command:

```
apt-get install <package>
```

Again, replace **<package>** with what you want to install, and you can use regular expressions. So, if you want to install all KDE-related packages, you'd use **apt-get install ^kde** – although that's probably not a smart move! Once you're done installing – really, APT is a shockingly easy way to install software – you can use this command to remove software:

```
apt-get remove <package>
```

MANAGE AUDIO FILES

Build your own MP3 jukebox

Maurice Kelly shows how to get the most from your MP3 collection using Linux.

The sad thing about Linux is that it isn't very often that people actually get to do something which hasn't been done before, so in many situations 'rolling your own' is pretty much an exercise in re-inventing the wheel. This article could be accused of doing the same, but I prefer to think of it as refining or customising the wheel – after all, a tractor wheel does not fit a sports car – somebody had to refine the design somewhere! By the end of this tutorial you will hopefully have the basis for the further development of your own MP3 jukebox. What you'll need includes:

- a sound card
- Apache Web server
- PHP
- Perl and modules MP3::Daemon, Audio::Play::MPG123, IO::Socket::UNIX, MP3::Info, Pod::Usage.

The sound card is self-explanatory – if you want to output the sound you're going to need it! *Apache* and *PHP* provide an easy-to-use interface to the jukebox. The hard work however is done by the excellent *MP3::Daemon* Perl module written by John Beppu. We're going to use the *mp3* program provided with this module as a wrapper to the command-line MP3 player *mpg123*. Without getting too boring about it, *mp3* is a control program for *mpg123* which runs as a background process churning out the tunes for you and I.

Hit the CPAN

Before we can go much further we're going to have to hit CPAN and get the required modules. If you're using a fairly recent distro

you'll most likely already have *IO::Socket::UNIX* and *Pod::Usage*. You might also have a nice and handy CPAN shell, but if not, it's not that hard to compile and install Perl modules from tarballs (which are supplied on the CD/DVD.) If you're using the CPAN shell for the first time, then you will want to **su** to root, and issue the following commands:

```
# perl -MCPAN -e shell
```

You'll be prompted for answers to lots of questions – for the most part the defaults are fine. If you want to change the answers

you can re-enter the CPAN shell at a later stage and issue the command **o conf init**. Once you're completely configured, you'll be left at the **cpan>** prompt where you should tell it to **install MP3::Daemon**. It will most likely tell you that you haven't got *Audio::Play::MPG123* installed, and ask you if you want to install it before *MP3::Daemon* – you should say yes for this and other pre-requisites. Once done you should see a number of messages regarding the building and installing of your new modules, and when that's finished you should be back at the **cpan>** prompt.

We're not quite done yet, as the standard version of *mpg123* is not ideal. Luckily *Audio::Play::MPG123* provides an appropriate version which you will have to build yourself. You may want to remove the existing *mpg123*, or copy it to a safe location. The following commands should create the modified version of *mpg123*:

```
cpan> look Audio::Play::MPG123
# cd mpg123
# make linux-help
# make linux-mmx
# make install
# exit
cpan> exit
```

(Note that where I performed **make linux-help** above, you will be shown a list of build types for various architectures – I chose to **make linux-mmx** as it was most appropriate for my hardware – you should choose what best suits your machine.)

At this point you should have a working installation of *MP3::Daemon* and the basis for the jukebox is ready. In fact, you could give up now as you have a reasonably powerful MP3 player on your system with an editable playlist, that you can control easily from the command line (hint: **mp3 help** or **man mp3** should give plenty of info on what you can achieve with this command.) But we're going to take this jukebox further by sticking on a web-based front end for more convenient control over the playlist (and, if you're in a networked environment, remote control over your jukebox.)

Control yourself


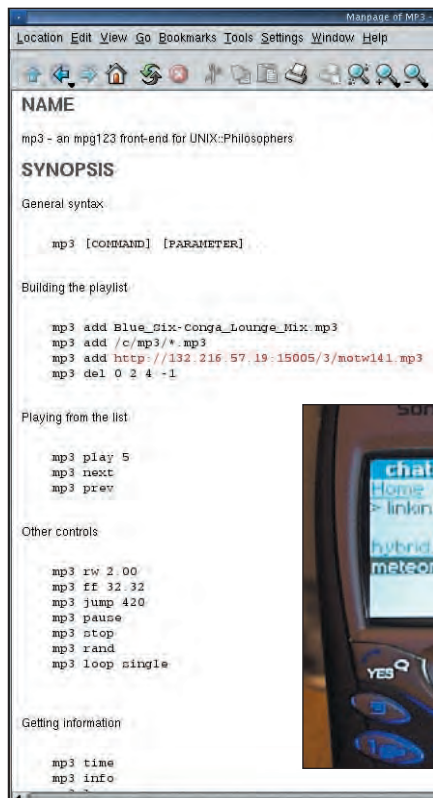
In order to control the *mp3* program via *Apache*, we need to make a very slight modification to it. Change to the root user and issue the following (assuming that *mp3* is installed in */usr/bin*):

```
# mv /usr/bin/mp3 /usr/bin/mp3jb
```

Now edit */usr/bin/mp3jb*, and change the **\$conf_dir** variable from **\$ENV{HOME}/.mp3** to **/tmp/.mp3** (on my copy of *mp3jb* this was at line 16.) The reason for this change is simple – I have found that when called by *Apache*, the **\$ENV{HOME}** variable does not appear to be set, and so the *.mp3* directory does not get created properly. Specifying a directory such as */tmp* (although you can choose anywhere you want really) means that the *.mp3* directory can be found. Of course it means that only one user can use that copy of *mp3jb* (hence our creating it as a separate copy of *mp3*.)

Perform a **ps aux | grep httpd** to determine what user your *Apache* process runs under, and then **su** to that user (on my Mandrake 9.0 installation, the user is called "apache.") You will probably have to **su** to root, then to the *Apache* user. At the

The *mp3* program is a basic implementation of the *MP3::Daemon*: Simple class. That said, it does have a lot of great features which you can find listed in the man page **man mp3**.



A WML-based interface to a jukebox – the format makes it less fully featured, but it's a lot more convenient when you really can't be bothered getting off that comfy couch!

command line try playing an MP3 file using `/usr/bin/mp3jb play /path/to/mp3_file.mp3` – hopefully you should hear something, but if not then you may have a permissions problem with your sound device – do an `ls -l` on `/dev/dsp` (or `/dev/sound/dsp` if your distro uses *devfs*.) If the permissions are wrong, change them to allow access to the *Apache* user (either by opening them up completely, or adding the user to a group which can access the device.) Note that some distros may have security features which routinely reset the permissions on certain files and devices to secure values, so you may need to update the scripts which do this.

You should next create a directory for your PHP control scripts to reside – you could create yourself a virtual host, or just a directory of your main DocumentRoot, it's completely up to you. Within the directory create an `index.php` file and put the following in it:

```
<html><head><title>MP3 Jukebox</title></head>
<body><pre>
<?
shell_exec( "/usr/bin/mp3jb play /path/to/mp3_file.mp3" );
$strOutput = shell_exec( "/usr/bin/mp3jb ls" );
echo( $strOutput );
?>
</pre></body></html>
```

Accessing the `index.php` page through your Web browser should kick off some audio output and generate a quick listing of the current playlist (most likely one song.) If you're hearing audio at this point, then the hard part is over, and it's time for your creative side to take over by writing an amazing Web-based interface (read on for more ideas.) If you're not hearing anything then you'll have to use a bit of patience to debug the setup – sadly Linux distros can vary greatly, so what works on one may need tweaking on another. Don't give up, and try the *LXF* forums for advice if you're stuck.

I'm a control freak!

You've got the 'infrastructure' in place – now all you have to do is tart it up a bit with some great looking Web pages. The example script above is all well and good, but isn't a lot of use for anyone wanting to play more than the song hard-coded into it. The simplest thing to do is modify the script so that it lists the available songs, and gives the option to add an individual song to the playlist.

The following snippet of code will display all the MP3s:

```
$strDirName = "/path/to/mp3files";
$intFileHandle = opendir( $strDirName );
$intFileKey = 0;
$strFileNameStore = "";
while ( $strFileNameStore = readdir( $intFileHandle ) )
{
    // Only include MP3 files
    if( preg_match( "\.mp3$", $strFileNameStore ) )
    {
        $arrFileEntry[ $intFileKey ] = $strFileNameStore;
        $intFileKey++;
    }
}
closedir( $intFileHandle );
sort( $arrFileEntry );
reset( $arrFileEntry );
while ( list( $key, $val ) = each( $arrFileEntry ) )
{
    if ( $arrFileEntry[ $key ] != "" )
    {
        $strShortPath = $arrFileEntry[ $key ];
```

Getting adventurous

The great thing about Linux is always being able to take things one step further. So now you've got yourself a nifty little MP3 jukebox – why not make it into the ultimate home audio system by piping the sound to a hi-fi, and grabbing yourself some sort of touch-screen monitor to act as an interface? Or why not create a WML interface to the system to allow control via a WAP gateway and a mobile phone? Obviously this requires a permanent Internet connection, a bit of WML

knowledge and a little *Apache* tinkering to allow PHP to serve digestible WML to WAP gateways.

Another possibility is to stream MP3s from your web-server. Many decent MP3 players are capable of playing MP3s across a network via HTTP. All you have to do is generate a Web page which lists the full URLs to your MP3s, set its Content-type to be `audio/x-mpegurl`, and tell your browser to associate that MIME type with your favourite MP3 player.



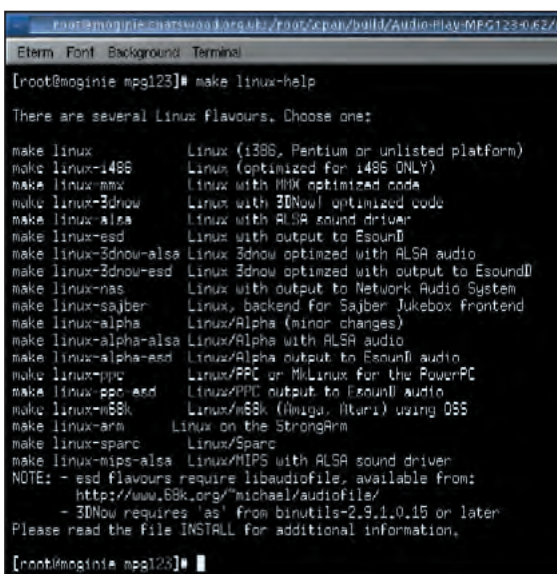
This article is aimed at the creation of a web-based MP3 jukebox, such as the one shown, but there's no need to limit yourself – it's just as easy to access the mp3 interface from a graphical toolkit such as QT.

```
echo( "<a href='\"index.php?\"' );
echo( \"subFileName=$strShortPath&\" );
echo( \"subAction=add\">\" );
echo( \"$strShortPath</a> \" );
echo( "<br>\\n\" );
}
```

You'll note that on line 24 we specify a **subAction** of **add** – we use this parameter to tell our control script that we want to add to the playlist the file specified by **subFileName**. The following snippet will do this for us:

```
if( $subAction == "add" )
{
    shell_exec( "/usr/bin/mp3jb add $strDirName/$subFileName" );
}
```

You now have a few ideas as to where you can take your own jukebox. Unfortunately we don't have enough space in the magazine to print a fuller control script, but have a look on the CD/DVD in the magazine section for an expansion on the framework above. [LXF](#)



The full list of Linux build options for the copy of *mpg123* that comes with *Audio::Play::MPG123*.

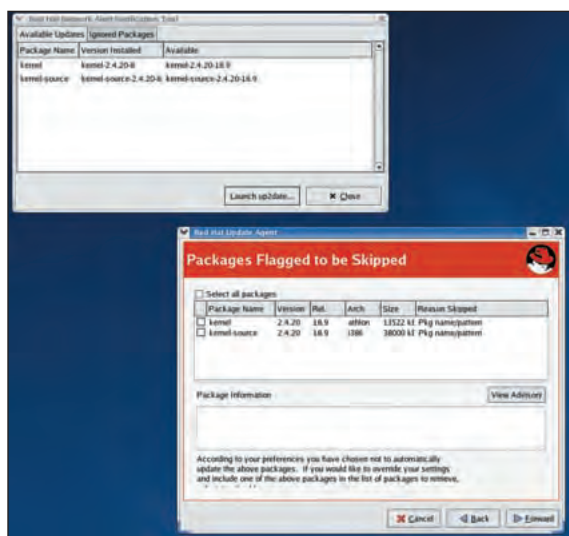


Fig2 The Red Hat Network will notify you of any available updates and that red button will pulse annoyingly if one is to be had. Here, all that are available are kernel updates that should never be installed via *up2date*. Had there been other updates, *up2date* would have downloaded and installed them for us had we told it to.

had to be an alternative way. We found several, but first let's quickly examine the underpinnings of them all, RPM itself.

RPM – the RPM Package Manager

One of the distinguishing characteristics of Red Hat is the way it assembles an application and manage it. RPM (only in more recent times) stands for RPM Package Manager and the knowledge about these packages and the RPM applications are codified in the book *Maximum RPM* (ISBN 1-8881-7278-9). You may read it to your heart's content at www.rpm.org/max-rpm/. Back then it was known as the Red Hat Package Manager, and its original purpose was to organise, distribute, install and manage software for a relatively new Linux distribution. It has since been ported to several commercial UNIX systems and is one of four predominant packaging systems used in Linux (.rpm, .deb, .tgz and source code tarballs). It is the packaging choice of all the major Linux distros except Debian (.deb packages) and Slackware (.tgz packages). No packaging system is perfect; each has its own strengths and weaknesses. Red Hat provides RPM because the development of RPM is inextricably linked to Red Hat itself, much as is GNOME and *ext3*.

An .rpm file provides a special binary header that tells us a great deal about the package: where it belongs, what files it contains, what other files it needs, and other information. The RPM application maintains a database of this information and provides a command-line interface to allow you to query installed and uninstalled packages, install them, remove them, and build them. RPM is designed for much more than simply installing and removing packages; it is a set of tools designed to create and manage packages built from pristine source for multiple architectures from a single source file, allowing for the use of GPG signatures for security and trustworthiness.

Obviously, RPM is a collection of totally unrepentant command-line tools. The main command itself offers some five dozen or so options, more than most anyone would likely care to master. It also allows you to install packages from remote ftp servers as well as from local drives or removable media.

Alternatives to Red Hat Up2date

Functionality for free

If you desire to use the client-server *up2date* service, there are two options to setting up your own servers to work with the existing Red Hat-supplied clients.

THE NHR-UP2DATE SERVER

This allows you to cache files locally and mimics the 'official' *up2date* server in its functions. Interestingly, they recommend that you establish your NHR-*up2date* server on a Red Hat 7.2 platform since that is what Red Hat uses for theirs (for now). It does require *perl-BerkeleyDB* that is not distributed with Red Hat, but is available from the site.

CURRENT

Current is another clone of the *up2date* server.

It doesn't appear at first glance to be as fully developed as NHR-*up2date*, especially since it suggests extensive testing before use and suggests suitability only for smaller networks, but this could be over-caution on the author's part. It does work on any recent version of Red Hat, however. The goal is to provide full *up2date* functionality; the installation documentation is very well written.

NOTE: Neither option outlined here provides you with a replacement for the Red Hat Packaging Tool, so you would need to use the command line RPM for managing custom or non-Red Hat RPM packages.

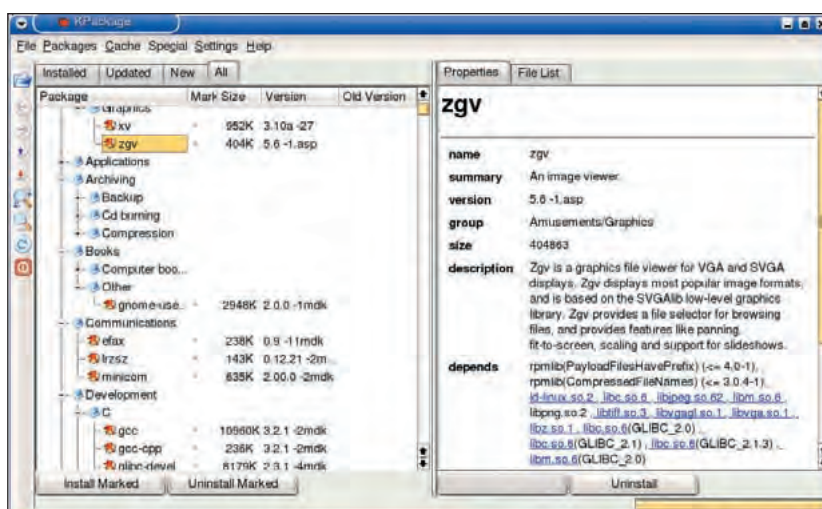


Fig3: Also gone but not forgotten is *Kpackage*, still available in the unexpurgated KDE.

You would think that as the trailblazer in RPM, Red Hat Linux would have a clever GUI package management tool. For a long time, Red Hat used *xrpm*, an application so outdated it no longer runs on modern Red Hat systems. The KDE desktop provides *kpackage* (Fig3) and the Gnome desktop provides *gnorpm* (Fig4) for package management, but neither of these graphical interfaces can fully exploit the power of the RPM system nor deal adequately with its biggest problem: Dependency Hell.

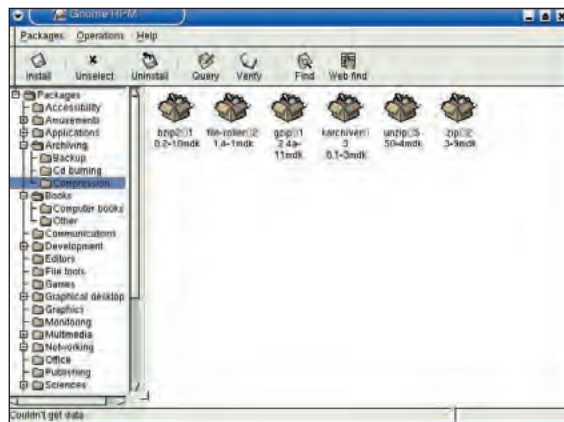
All of the dependency information is contained in the RPM files themselves, so it would seem to be a simple matter to maintain a database of the dependencies and let the computer sort it out. That's just what someone thought to do, not once, but as is common in the Open Source/Free Software world, several people thought to do it in slightly different ways bringing us the *apt*, *urpmi* and *yum* applications, all of which are elaborate, "smarter" wrappers to RPM itself; they all use RPM as the core and wrap additional functionality around it. (It's these database files that make APT, URPMI, YUM and other applications large and unwieldy. The .deb packages handle dependencies in a different fashion from RPM and so need smaller database files; however, they lack the ability to differentiate between a package optimised for i386 versus i686 as .rpm files can.)

While there are more alternatives than will be discussed here, this tutorial will focus on three applications: APT, URPMI, and YUM. All operate in a similar manner and all attempt to address



TutorialPackageManagement

Fig4: Familiar to Gnomistas worldwide, *gnorpm* has fallen from grace at Red Hat.



◀ the same set of problems. What you choose is up to you, but we'll first offer up our favourite.

APT

As we mentioned earlier, the Debian distribution (and other distributions built upon it like Xandros) use the Debian Packaging System, built around a program named *dpkg*, which is similar to RPM. A group of command line tools collectively named APT (Advanced Package Tool) was written to manage *.deb* packages. It was only matter of time until someone hacked APT to allow it to handle *.rpm* packages as well as *.deb* packages. The first major distribution to do so was Connectiva. The Connectiva team also developed a GUI front-end for APT originally called *Raptor*, but now known as *Synaptic*.

Fortunately for us, both APT and *Synaptic* are available in Red Hat-compatible *.rpm* files. Let's install them straightaway, configure them, and see how they work. There are two main sites for obtaining the packages as well as custom packages for Red Hat: FreshRPMs and Fedora. Either is acceptable, but we'll go

with FreshRPMs for our example; you should explore both options. The weblinks for both follow in the body of this tutorial.

Beginning with APT/Synaptic

In order to use *apt-get/Synaptic* effectively, you need to:

- Install APT
- Synchronise the database
- Install *Synaptic* using the **apt-get** command

All three steps are easily accomplished. Before we start, we should import the GPG key from FreshRPMs and let RPM know we have it. RPM can use the key to make certain that the packages we get are signed with the FreshRPMs PGP key.

Obtain an APT .rpm package built for Red Hat 9.0 and install it

The FreshRPMs.net site provides binary packages for Red Hat 6.2 through the current version. You can find a link to all the available versions at <http://apt.freshrpms.net/>. We'll use the remote access functionality of the *rpm* command to download and install APT. Log in as root and type (check the website first for an updated version or this won't work):

```
# rpm -ivh http://ftp.freshrpms.net/pub/freshrpms/redhat/9/apt/apt-0.5.5cnc6-fr2.i386.rpm
```

This will use your Internet connection to contact the FreshRPMs.net site, download the apt *.rpm* file and install it.

Synchronise the APT database file

```
# apt-get update
```

Which shows us:

```
Get:1 http://ayo.freshrpms.net redhat/9/i386 release [1169B]
```

```
Fetchd 1169B in 0s (1520B/s)
```

```
Get:1 http://ayo.freshrpms.net redhat/9/i386/os pkglist [1356kB]
```

```
Get:2 http://ayo.freshrpms.net redhat/9/i386/os release [140B]
```

```
Get:3 http://ayo.freshrpms.net redhat/9/i386/updates pkglist [225kB]
```

Dependency Hell

How to sort out your inter-related applications

Many applications require that other applications be installed on the system before they will work properly. As an example, the application *xrpm* requires that Python, *tix*, and *tkinter* be installed before *xrpm* can be installed and run, but we don't know that yet. Python's already been installed on the system, so here's what you'll see when you download *xrpm* and try to install it:

```
# rpm -ivh xrpm-2.2-8mdk.i586.rpm
```

error: failed dependencies:

```
tix is needed by xrpm-2.2-8mdk
```

```
tkinter is needed by xrpm-2.2-8mdk
```

What are those applications? Where can they be found? What version do we need? How must they be compiled and configured? The application *xrpm* depends on these other applications. How are you to easily know? Where are these packages found? In which order must they be installed? Where is this information kept? Argh! It would be even worse if each of those two apps depended on others and so on. Welcome to Dependency Hell.

Here's how the *urpmi* application handles the same installation (when we have properly configured repositories):

```
# urpmi xrpm-2.2-8mdk.i586.rpm
```

To satisfy dependencies, the following packages are

going to be installed (3 MB):

```
tix-8.3.3-21mdk.i586
```

```
tkinter-2.2.2-6mdk.i586
```

```
xrpm-2.2-8mdk.i586
```

Is this OK? (Y/n) y

```
ftp://ftp.my.site/pub/mandrake/9.1/9.1-tree/i586/Mandrake/RPMS/tix-8.3.3-21mdk.i586.rpm
```

```
ftp://ftp.my.site/pub/mandrake/9.1/9.1-tree/i586/Mandrake/RPMS/tkinter-2.2.2-6mdk.i586.rpm
```

```
installing /var/cache/urpmi/rpms/tix-8.3.3-21mdk.i586.rpm /var/cache/urpmi/rpms/tkinter-2.2.2-6mdk.i586.rpm xrpm-2.2-8mdk.i586.rpm
```

```
Preparing... #####
```

```
1:tix #####
```

```
2:tkinter #####
```

```
unable to access rpm file [xrpm-2.2-8mdk.i586.rpm]
```

```
error: read failed: Success (0)
```

```
# urpmi xrpm-2.2-8mdk.i586.rpm
```

```
installing xrpm-2.2-8mdk.i586.rpm
```

```
Preparing... #####
```

```
1:xrpm #####
```

Note that the install failed at the end because *xrpm-2.2-8mdk.i586.rpm* was just a locally downloaded file,

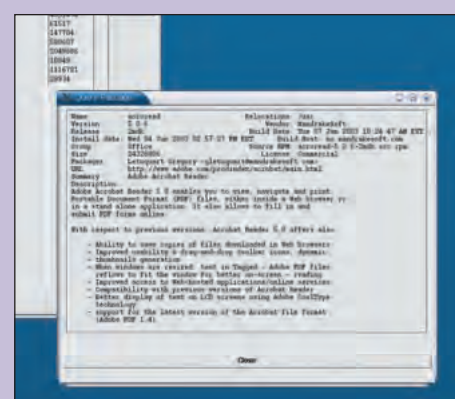


Fig5 A friend to many a Red Hatter, *xrpm* is no longer provided by Red Hat in order to focus development on the Red Hat Packaging Tool. With a little effort, you get that old, warm retro feeling.

not a file known to the *urpmi* database. Once the dependencies were met, it installed with ease. *apt* works essentially the same way. The results are shown above in fig5.


```
Get:4 http://ayo.freshrpms.net redhat/9/i386/updates release [153B]
Get:5 http://ayo.freshrpms.net redhat/9/i386/freshrpms pkglist [135kB]
```

```
Get:6 http://ayo.freshrpms.net redhat/9/i386/freshrpms release [157B]
```

```
Fetchd 1716kB in 34s (49.2kB/s)
```

```
Reading Package Lists... Done
```

```
Building Dependency Tree... Done
```

APT looks at the repository specified in the `/etc/apt/sources.list` file, the contents of which are provided by FreshRPMs already configured for their repositories for Red Hat 9; other versions are listed and commented out. Here's a sample of the file content:

```
# Red Hat Linux 9
```

```
rpm http://ayo.freshrpms.net redhat/9/i386 os updates
freshrpms
```

```
rpm-src http://ayo.freshrpms.net redhat/9/i386 os updates
freshrpms
```

In this example, the rpm repository at psyche.freshrpms.net is the source. The format is described in the man page for `sources.list`. The format is: **rpm URI component1 component2 component3** etc.

The Universal Resource Indicator points to the base directory of the repository. Each component serves to differentiate the location of groups of files within that repository. This structure is a by-product of how Debian organises ftp file trees.

The URI can be a local or network directory, a CDROM, an ftp or http server and can employ `rsh` or `ssh` if necessary. The man page, fortunately, provides extensive examples.

Since every source must be synchronised and kept current, the command `update` command is the first you use. It should not be necessary to use it again on directories where the files will not change (the base directory, for example).

Install Synaptic using apt-get

It's as easy as:

```
# apt-get install synaptic
```

which shows us:

```
Reading Package Lists... Done
```

```
Building Dependency Tree... Done
```

```
The following NEW packages will be installed:
```

```
synaptic
```

```
0 packages upgraded, 1 newly installed, 0 removed and 36
not upgraded.
```

```
Need to get 376kB of archives.
```

```
After unpacking 1159kB of additional disk space will be used.
```

```
Get:1 http://ayo.freshrpms.net redhat/9/i386/freshrpms
synaptic 0.36.1-fr1 [376kB]
```

```
Fetchd 376kB in 7s (47.8kB/s)
```

```
Executing RPM (-Uvh)...
```

```
Preparing... ##### [100%]
```

```
1:synaptic ##### [100%]
```

That's all there is to it. Once Synaptic is installed, launch it with `# synaptic &`

You may also use the desktop menu entry for *Synaptic*, or manually create an icon for your desktop. No matter how you start it, you'll see a GUI appear that looks somewhat similar to those of *kpackage* and *gnorpm*. The biggest improvement over those applications, however, is that APT/*Synaptic* will handle all the dependencies for you.

Command Line or GUI?

The ability to use APT from the command line can be valuable if you want to script its behaviour (for regular updates via *cron*) or will be using it on a server with no X11 display. The *Synaptic* GUI can make the job easier if you are using a graphical desktop as it more conveniently displays packages that are available for upgrading and installation. Let's examine both the command-line and GUI tools.

Using APT at the Command Line

The command line interface to APT is the command **apt-get instruction**. With the **apt-get** command you can use one of the following instructions:

- **update** – synchronises the local package index with that of the repository.
- **upgrade nameofpackage** – installs the newest version of the package.
- **dist-upgrade** – upgrades your entire system to the latest versions available from the repository.
- **install nameofpackage** – install the package and all dependencies.
- **remove nameofpackage** – erases (removes or uninstalls) the package.

The man page contains references to a number of useful command options like **-y** to assume a "yes" answer to all queries, **-s** to simulate an installation, and **-u** to display a list of packages to be upgraded among others.

If you are installing a single package and know exactly what you want to do, the command-line power of **apt-get** is seductively simple and easy. And of course, **dist-upgrade** is a powerful command and the pride of the Debianistas. As an example, running it on our test system showed us:

```
# apt-get dist-upgrade
```

```
Reading Package Lists... Done
```

```
Building Dependency Tree... Done
```

```
Calculating Upgrade... Done
```

```
The following packages will be upgraded:
```

```
cups cups-libs eog ethereal evolution foomatic gaim
ghostscript glibc glibc-common glibc-devel gnupg grip gthumb
```

Debian Packages

Debian packages end with a `.deb` extension and are used for the Debian distro and Debian-derived distros. The `.deb` package is simply two tarballs wrapped in a GNU archive. Very similar to RPM in functionality, the policy that governs all aspects of `.deb` files can be found at this address: www.debian.org/doc/debian-policy/

Using GPG Keys with RPM

RPM packages can be signed using GPG. Red Hat signs all its packages and many packagers sign theirs. You'll typically find the public PGP key with the files on the website. Download it and install it for FreshRPMs.net with:

```
# rpm --import RPM-GPG-KEY.txt
```

It's not mandatory to install it, but RPM will warn you if you attempt to install a signed package without the key; ignore the warning or install the key.



Alternative Alternatives and Mixing Repositories

In addition to FreshRPMs, the Fedora project offers a slightly different approach to providing **apt-gettable** packages for Red Hat. Fedora is a group effort with elaborate guidelines for participation. Fedora wisely recommends against "mixing" repositories (using both FreshRPMs and Fedora, for example) because of potential package version conflicts (a very

real and annoying circumstance). Each project has decided to follow slightly different versioning policies and their packages for the same app can, well, give poor little APT major indigestion. Be careful which repositories you add to your list. As always with Linux, carefully examine all your options to find a solution that fits your needs.

TutorialPackageManagement

gtkhtml hpijs httpd httpd-manual kbase kdegraphics kdelibs krb5-devel krb5-libs mod_ssl nmap nsd openssl openssl-devel samba samba-client samba-common sendmail sendmail-cf tcpdump xinetd xpdf

The following NEW packages will be installed:

gtksPELL id3lib libexif

36 packages upgraded, 3 newly installed, 0 removed and 0 not upgraded.

Need to get 106MB of archives.

After unpacking 2962kB of additional disk space will be used.

Do you want to continue? [Y/n]

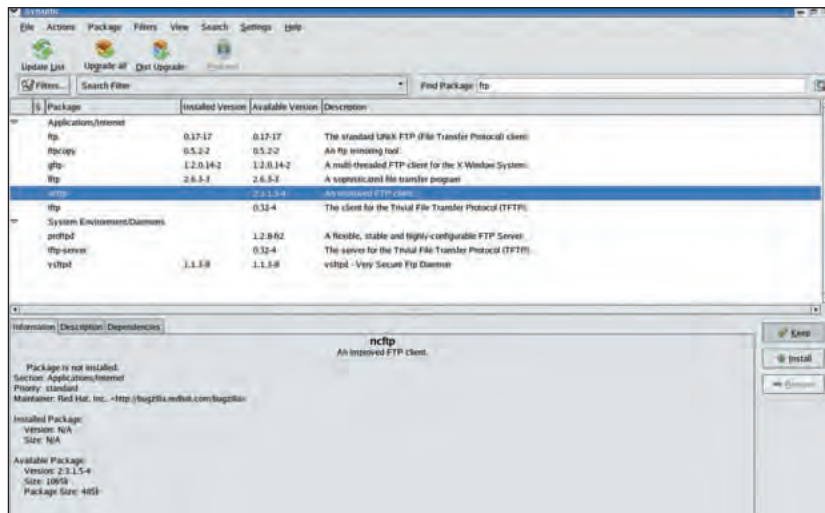
After we replied 'Y', the files were downloaded, installed and our system upgraded as FreshRPMs.net maintains all the current files Red Hat provides.

The APT/Synaptic Graphical Interface

Synaptic is just a pretty GUI front-end to APT. Some information is better viewed with a graphical presentation and some people simply wish to use a 'point-and-click' style interface; you have a choice. The main screen is shown in **Fig6** below.

Beginning in the upper left, the Update List button is the equivalent to **apt-get update** which syncs the file database. Upgrade All will perform an **apt-get upgrade** and Dist Upgrade is the equivalent to **apt-get dist-upgrade**. The middle window provides a file listing which can be filtered in a variety of ways. When the packages are displayed, they will show the installed

Fig 6: The Synaptic interface does everything from one screen, unlike URPMI which uses four! Here, we've filtered on any package with 'ftp' in the name to see what we have available. The ncftp app is selected and information about it appears at the lower left. All we need do is click on 'Install' and then 'Proceed'.



Comparison of Selected Features

The package management contenders go toe-to-toe

Feature	APT	URPMI	YUM
Create file index from local repository without extra tools	N	Y	Y
Create file index from remote repository without extra tools	N	N	Y
Advanced command line database queries	N	Y	N
GUI available	Y	Y	N
Easily installed	Y	Y	N
Easy database synchronisation	Y	Y	Y
Available for earlier versions of Red Hat	Y	N	Y
Small codebase	N	N	Y
Legacy (.deb) free	N	Y	Y
GPL	Y	Y	Y
Well documented	Y	N	Y
Compatible with RHN/up2date	N	N	N

and available version and a brief description. The lower left area provides detailed information; buttons on the left trigger actions. You may even edit the `/etc/apt/sources.list` from the Settings menu, should you wish to..

URPMI

The Mandrake distribution was initially a massaged version of Red Hat with the KDE desktop as the default. It became apparent to Mandrake that `xrpm/kpackage/gnorpm` were inadequate to the task and they set about building URPMI, an RPM frontend of their own patterned after APT. Why they did so is unclear, but perhaps they viewed Connectiva's approach as a kludge. More likely, since the new Mandrake GUI installer was written in Perl, it made sense to write URPMI in Perl as well (APT is written in C++).

Why has URPMI not spread to other distros since it is licensed under the GPL? The GUI is tied closely to Mandrake, and documentation for creating the files necessary for a repository has not been widely available. Just because it's GPL doesn't mean they have to make it easy for you!

The URPMI cadre of applications fall into three categories: Repository Management, Package Management, and Database Query. We will examine these below in a basic manner.

Repository Management

Repositories can be added or removed with the `urpmi.addmedia` and `urpmi.removemedi` commands. The `urpmi.update` command downloads a current file list (named `hdlist.cz`) from the repositories which must have such a file list pre-existing; all official Mandrake mirrors have such files.

The syntax for adding an ftp repository is:

```
# urpmi.addmedia nameofrepository
ftp://domain/and/path/to/Mandrake/RPMS with ../base/hdlist.cz
```

The name can be anything you choose; a list of official mirrors can be found at www.mandrakesecure.net/en/ftp.php; the part **with** `../base/hdlist.cz` is necessary to point URPMI to the adjacent directory where the files actually reside.

If you are unsure about how to actually do this, an automatic repository generator, the *Easy URPMI Configurator*, can be found at <http://plf.zarb.org/~nanardon/urpmiweb.php>. Just select the options appropriate to you and click on the Generate button; copy and paste the provided configurations in a terminal window as the root operator.

The syntax for deleting a repository is:

```
# urpmi.removemedi nameofrepository
```

To obtain the most current index file listing:

```
# urpmi.update nameofrepository
```

to update one repository, or

```
# urpmi.update -a
```

to update all the repositories you use. It is not necessary to update the base and contribs repositories, because these are not changed once the version is released.

Package Management

You would use `urpmi` to (i)nstall or `urpme` to (e)rase a specific package. If there are dependencies, you will be provided with a list and given the option to install or delete them. You may also use a partial package name and `urpmi` will provide a list of possible candidates.

Database Queries

The `urpmq` command facilitates queries to the `urpmi` database. The command is

```
# urpmiq option
```


Resources

The RPM Homepage

www.rpm.org

The centre of the RPM universe.

The Maximum RPM book

www.redhat.com/docs/books/max-rpm/index.html

The canonical reference, outdated but “under construction”.

The RPM HOWTO

<http://www.rpm.org/RPM-HOWTO/>

A *tour-de-force* of what to do.

‘Is it time to change RPM?’

<http://freshmeat.net/articles/view/182/>

A discussion about the shortcomings of RPM files, essentially a suggestion that RPM developers adopt the Debian packaging policies.

Text or Console Apps

<http://rpm.redhat.com/software/console/>

A partial list of alternative applications; GRAB looks promising as it works with any ftp/http repository without requiring special tools or remote index files.

RPM Updaters

<http://rpm.redhat.com/software/updaters/>

Some more sophisticated applications including *up2date* replacements.

RPM GUIs

<http://rpm.redhat.com/software/gui/>

Pretty, clicky interfaces.

RPM Packagers

<http://rpm.redhat.com/software/packagers/>

A few interesting utilities, especially *checkinstall*, previously covered in *LXF* which will make an .rpm out of just about any source code so RPM can manage it.

Miscellaneous RPM utilities

<http://rpm.redhat.com/software/misc/>

RPM developers will find useful tools here as well as .tgz and .deb to .rpm conversion utilities.

Synaptic Homepage

<http://www.nongnu.org/synaptic/>

URPMI Homepage

www.linux-mandrake.com/cooker/urpmi.html

A Spartan page from the developers with broken download links. Talk about obscurity...

The URPMI HOWTO

<http://myweb.tiscali.co.uk/jwrobinson/docs/urpmi-howto/>

A non-official overview of URPMI created by a user with some helpful examples of how the components of URPMI are used.

UsingUrpmi

<http://mandrake.vmlinux.ca/bin/view/Main/UsingUrpmi>

Finally, a little official documentation.

Urpmi Mini-mini-howto

[www.pclinuxonline.com/modules.php?name=News&file=article&sid=3460](http://pclinuxonline.com/modules.php?name=News&file=article&sid=3460)

And some more useful user-supplied documentation.

apt+rpm HOWTO

<http://bazar.conectiva.com.br/~godoy/apt-howto/index.html>

A HOWTO from the folks who made PAT work with RPM. Provides a good discussion on using GPG signatures.

APT-enabled CDROMS

<http://freshrpms.net/apt/>

At the bottom of the page, there is an explanation on how to re-burn the Red Hat CDROMS to enable them for APT, useful if you need to use that media instead of a LAN or Internet connection.

Repositories Galore

There are several APT-enabled repositories with some interesting files. Check out the lists at

<http://freshrpms.net/apt/repositories.html>, for

multiple sources of Red Hat 6.x, 7.x, 8.x and 9 repositories as well as speciality packages. Also see the list of Red Hat RPM repositories available at

www.pclinuxonline.com

Some useful options are:

- **list-media** – lists all available repositories
- **headers nameofpackage** – extracts the package headers to *stdout*
- **d nameofpackage** – queries package dependencies
- **y partial nameofpackage** – use a “fuzzy” search on a partial name
- **R nameofpackage** – discloses what packages requires (or depends on) this package

Another query command is **urpmf**, which helps you find the package that contains individual files. For example, you’re trying to compile an application from source and you are told that it requires *jpegint.h*. What package provides that?

```
# urpmf jpegint.h
```

which replies:

```
mozilla-devel:/usr/include/mozilla-1.3.1/jpeg/jpegint.h
```

and we deduce that the package **mozilla-devel** contains the file we seek. Install it and continue. You can also search the package descriptions. Suppose we want to install a graphics application for X11, but aren’t sure what is available.

```
# urpmf -summary graphic -a X
```

which replies :

```
rox-session:Session manager for the ROX graphical desktop
```

```
rox-system:System monitor for the ROX graphical desktop
```

```
rox-wallpaper:Set the wallpaper image for the ROX graphical desktop
```

```
links-graphic:Lynx-like text/X11 WWW browser
```

```
xfig:An X Window System tool for drawing basic vector graphics
```

```
xboard:An X Window System graphical chessboard
```

```
xfishtank:An X Window System graphic display of an animated aquarium
```




Fig7 These two screens are the graphical interface to selecting and configuring URPMI repositories. Sadly, the configuration screen works less reliably for us than the command line **urpmi.addmedia** command.

xxgdb:An X Window System graphical interface for the GNU gdb debugger

The **-a** is a binary AND operator and **-o** is a binary OR operator; more options are explained in the man page. The **urpmf** command is more powerful than any provided with APT.

The URPMI Graphical Interface

Mandrake provides multiple applications as graphical front-end to the *urpmi* functions, each serving a different function. There is:

- **/usr/sbin/edit-urpm-sources.pl** to configure repositories,
- **/usr/sbin/rpmdrake-remove** to remove software,
- **/usr/sbin/rpmdrake** to install software
- **/usr/sbin/MandrakeUpdate** to install updates.



TutorialPackageManagement

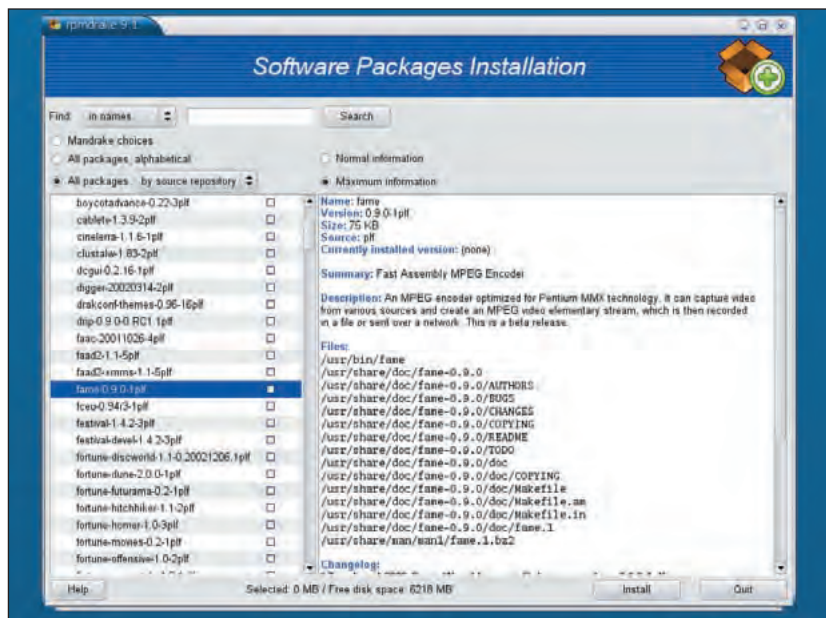


Fig8 You can search by keyword through several different indexes to locate the package you want. Dependencies are automatically selected and a click of a button installs the lot.

They are all accessed through the *Mandrake Control Center* (/usr/bin/mcc), or they can be called directly (although they are not intended to do so). The screen to configure and manage repositories is shown in Fig7 on the previous page, and the screen to search and install them is shown in Fig8 above.

While URPMI itself should be easy to port to Red Hat, the GUI would be difficult (at least for this author, who is most definitely not a programmer). There are no *urpmi* packages for Red Hat, but since *urpmi* is Perl code, it's possibly trivial to port it to Red Hat; the GUI tools may be more problematic as they rely on several Mandrake-specific packages. *urpmi* is examined in this article primarily so you could have a better understanding of the different approaches to advanced package management. You may want to port it to a different distribution; this would be possible since *urpmi* is licensed under the GPL.

YUM

APT for RPMs is seen by some as somewhat of a hack and an application with bloated code (remember that it was originally written to use .deb, and that .rpm support was added after the fact); the URPMI application was intended primarily for Mandrake. What other option might there be?

Enter YUM, written in Python (the Red Hat installer, *Anaconda*, is written in Python and some code is shared between the two). YUM is based on YUP, the Yellow Dog Updater provided with Yellow Dog Linux for the PPC (www.yellowdoglinux.com/). YUM packages are available for Red Hat 7.3, 8.0 and 9 and can be obtained through the YUM homepage at <http://linux.duke.edu/projects/yum/>. While a source tarball is provided, YUM must be installed as an RPM package. If you are a dyed-in-the-wool build-it-yourself type, you can build an RPM package from the tarball with `rpmbuild -ta yum-1.0.1.tar.gz`, since the tarball contains the .spec file necessary for package building. You'll find the new RPM file, `yum-1.0.1-1_80.noarch.rpm`, in `/usr/src/RPM/RPMS/noarch/`; install it with

```
# rpm -ivh /usr/src/RPM/RPMS/noarch/yum-1.0.1-1_80.noarch.rpm.
```

Of course to build the .rpm package, you must have *rpmbuild* and all its dependencies installed along with any other dependencies from *rpmbuild* – have we mentioned Dependency Hell before?

YUM consists of two applications, *yum* and *tum-arch*.

- **/usr/sbin/yum-arch** – run this command in the directory of the repository, this application generates the database that *yum* uses.
- **/usr/sbin/yum option** – the main application; you will use this with an appropriate option to make YUM work, as follows:
- **check-update** – tells you if any updated packages are available.
- **Update** – updates all installed packages if updates are available.
- **install nameofpackage** – install the requested package and any dependencies.
- **remove nameofpackage** – uninstalls the named package.
- **List** – list all packages in the repository that can be installed.
- **Clean** – by default, downloaded packages are not deleted after installation, so you need to clean up the RPM files manually.

More options are discussed in the man page for YUM. As of yet, there is no GUI for YUM and it remains solely a command-line application. YUM is licensed under the GPL. It is a good candidate for porting to other distros. Once YUM is installed but before you can run it, you need to edit the file `/etc/yum.conf` to list a repository. The default file looks like this:

```
[main]
cachedir=/var/cache/yum
debuglevel=2
logfile=/var/log/yum.log
pkgtpolicy=newest

[base]
name=Red Hat Linux $releasever base
baseurl=http://mirror.dulug.duke.edu/pub/yum-repository/redhat/$releasever/$basearch/

[updates]
name=Red Hat Linux $releasever updates
baseurl=http://mirror.dulug.duke.edu/pub/yum-repository/redhat/updates/$releasever/
```

Note that the default file points to the repository at Duke University in Durham, North Carolina, USA (no surprise since Duke University holds the copyright on YUM). An example entry for `rpmfind.net` (taken from <http://rpmfind.net/linux/freshrpms/ayo/>) would be:

```
[os]
name=Red Hat Linux 9 os
baseurl=http://rpmfind.net/redhat/9/i386/os

An example entry for the Fedora project (www.fedora.us) is:

[fedora-us-9]
name=Red Hat Linux 9 -- Fedora US mirror
baseurl=http://SERVERNAME/fedora/redhat/9/i386/yum/os/

[fedora-us-9-updates]
name=Red Hat Linux 9 updates -- Fedora US mirror
baseurl=http://SERVERNAME/fedora/redhat/9/i386/yum/updates/

[fedora-us-9-stable]
name=Fedora Linux (stable) for Red Hat Linux 9 -- Fedora US mirror
baseurl=http://SERVERNAME/fedora/redhat/9/i386/yum/stable/
```

You need to replace **SERVERNAME** with a server's name from Fedora's list found at www.fedora.us/wiki/FedoraMirrorList.

The files from FreshRPMs and Fedora can be used to add much multimedia functionality that Red Hat has chosen to remove or not provide, like *Mplayer* (for DVD viewing) and MP3 support. They're all just an **apt-get install** away. That should keep you busy until the next issue of LXF arrives.

NEXT MONTH

We examine how to mirror a repository locally using tools like *ftpcopy*, *curl*, *rsync*, and *fmirror* and configure *apt*, *urpmi*, and *yum* to use our local repositories. As a bonus, we can use these local repositories to effect ftp installs of Red Hat on our local LAN.

Until then, since you are APT-enabled, you have a few interesting options: If you want a 'stock' KDE 3 desktop for your Red Hat, the KDE for Red Hat Project has APT files available; you'll be amazed at just how easy a KDE installation can be when using APT. It's at <http://kde-redhat.sourceforge.net/>. Even Lindows users aren't left out of APT at www.linuxdownload.com/ and it's FREE!

LINUX/WINDOWS INTEGRATION

Authentication and Authorisation in Samba

PART 2 Authentication and authorisation are the Tweedledum and Tweedledee of access control. They're separate, but you never see them apart, says **Dr Chris Brown**.

Actually, it's hard to over-estimate the importance of *Samba* in facilitating Linux / MS Windows integration and the adoption of Linux into the commercial marketplace. The quality of the work being done by the *Samba* development team certainly deserves our recognition and support. Last month, we began our tour of *Samba* by looking at the overall picture of file sharing in Windows. We talked about workgroups and shares, and about the SMB/CIFS protocol that makes it all work. We saw that *Samba* is a suite of programs and services that can provide file and print sharing services to Windows clients, and took a brief look at its configuration file, `smb.conf`. We learned about *Samba*'s web-based graphical configuration tool, SWAT (*Samba* Web Administration Tool). We also saw that *Samba* has client-side components that allow you, for example, to mount shares exported by Windows file servers onto a local Linux directory.

This month, we'll complete the tour by looking at the business of authentication and authorisation in *Samba*. Let's start by making sure we know what 'authentication' and 'authorisation' mean.

Authentication is the business of making certain that you know the identity of a user who is requesting a service – in this case, the service in question is access to a file share or a printer exported by a *Samba* server. From the end user's viewpoint, authentication is almost always just a matter of supplying a username and a password. Authorisation is the business of deciding whether a specific user is allowed to access a specific resource; for example, is user jane allowed to access the share she's trying to connect to, and if so, is she allowed to read and write the files in that share?

Authentication and authorisation usually go together. Authorisation checks against a user identity are meaningless unless you're sure you know who the user is. Conversely, authentication is a bit pointless unless you're subsequently going to use your knowledge of the user's identity to control what he can do. Nonetheless, it's important to remember that authentication and authorisation are actually separate operations. *Samba* always needs to have a Linux UID to run with, corresponding to the user on the windows client.

Behind the scenes

It will help us understand *Samba*'s security mechanisms if we know a little bit about how it works behind the scenes. When *Samba*'s `smb` server starts up, there is initially just a single parent process running as root. When a user on a Windows machine connects to a *Samba* share, a new TCP connection is opened to the *Samba* server. The server will authenticate the user and check the user identity against the access control directives in its

configuration file. If the user is allowed to connect to the share, *Samba* starts up a child process to service that connection. When the client accesses a file in the share, the child process switches to run under the identity of the client user. Requests to read and write any files within the share will then be checked against the usual 'rwx' file permissions for that user. These checks aren't made by *Samba*, they're just Linux's underlying security model being applied. The overall process, including some details we'll get to in a minute, is shown in **Fig1** on the right.

There are a couple of points to note about all this. First of all, access controls are applied at two levels; first by *Samba*, as determined by entries in its own config files, and second, by the underlying Linux filesystem. An important consequence of this is that *Samba* needs to establish an identity for the user which is known to Linux, corresponding either to an entry in `/etc/passwd`, or the NIS `passwd` map, or whatever user account database the system is configured to use. To put it in its most simple terms, *Samba* always needs to have a numeric Linux user ID (UID) corresponding to the client user.

Why is this an issue? With NFS-based file sharing (to digress for a moment) Linux simply assumes that the UID of the user on the server is the same as his UID on the client. So if user mike on a Linux client has the UID 507, NFS uses the UID 507 on the server, too. In the case of SMB-based file sharing, the client machine is probably some version of Windows, and of course, Windows simply doesn't use UNIX UIDs. So in this case, it's the user name that's transferred to the server – hence the need to be able to resolve the name onto a UID.

Encrypted passwords

Now a bit of history. Early versions of Windows, such as Windows 95, didn't really know about users at all. You didn't have to log in on the machine. It's true that if the machine was configured to use Microsoft networking, you'd be asked for a username and password when the machine booted. But it didn't use that information to decide whether you were allowed to use the machine, it simply squirrelled it away to use later if the machine tried to connect to an SMB share. Then it simply passed the name and password, in cleartext, to the server. In the case of a *Samba* server, it could use these to authenticate against a regular UNIX account, thus establishing a UID for the client.

Later versions of Windows, such as NT, 2000, and XP, mandate a proper authenticated login to the machine. More importantly, Microsoft changed the authentication scheme in SMB to use a challenge-response mechanism based on an encrypted password. Of course, the encryption method was not compatible with the one used in UNIX. Whilst this was a good step forward for

Author info

Dr Chris Brown is a UNIX and Linux consultant, trainer and writer of many years standing, though he claims to have spent most of them sitting down. He recently added an RHCE to his list of qualifications and reckons that, apart from taking his car for its MOT, it was the most stressful thing he's done for decades. Chris can be contacted at: chris@idlearn.com

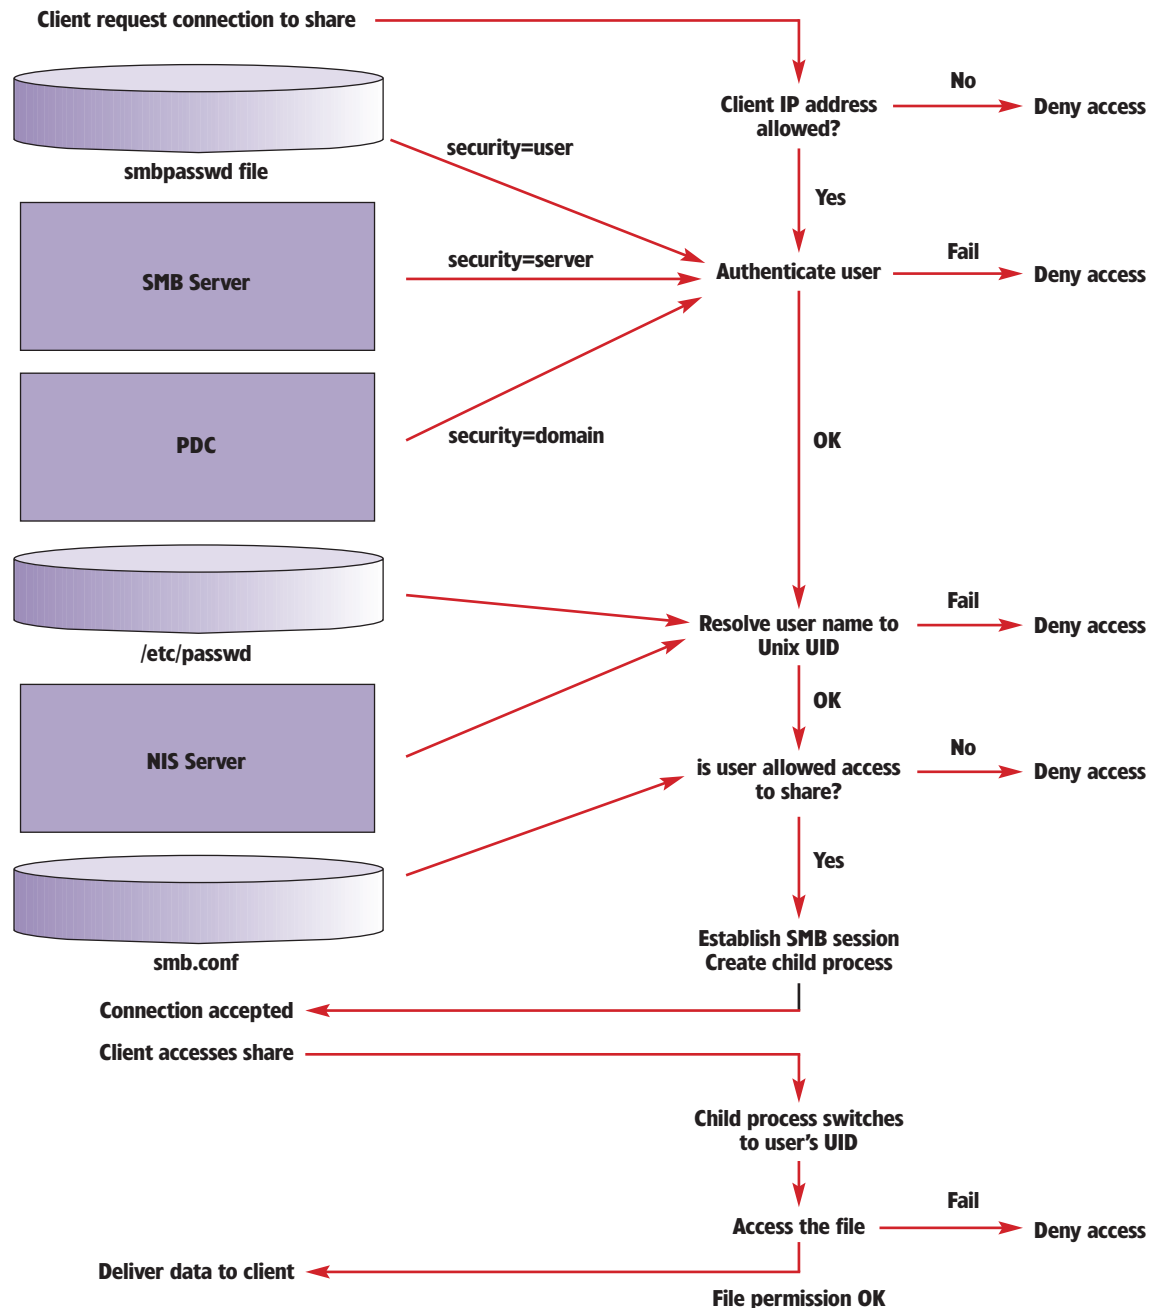


Fig1 An overview of how *Samba* works behind the scenes in making decisions to control access.

Windows, it caused consternation in the *Samba* world because the server no longer received a cleartext password and was thus not able to authenticate against the UNIX account database.

Samba can still be made to use cleartext passwords by including the directive

```
encrypted passwords = no
```

in the `smb.conf` file. It even provides patch files to doctor the registry on some versions of Windows to revert to the use of cleartext passwords. But cleartext passwords are not really a good idea, and this approach is, as they say, deprecated.

The Samba Password Database

To cope with the encrypted password mechanism, *Samba* was forced to use its own account database in which the windows user names and encrypted passwords could be stored. On my Red Hat system, this file is `/etc/samba/smbpasswd`, though if you've installed *Samba* from source, it might be elsewhere. There's also a program called `smbpasswd` that is provided as part

of the *Samba* suite to manage this file. This program can be invoked by a normal user to change their *Samba* password, in much the same way that the `passwd` command is used to change regular Linux passwords.

The superuser can run `smbpasswd` to change any user's password or to add new users to the file. For example, the command:

```
# smbpasswd -a kim
```

will add the user `kim` to `smbpasswd`. It will also prompt for, and set, `kim`'s password. For this command to succeed, `kim` must already have a regular Linux account.

The four security modes

To round off our discussion of authentication we should discuss the four 'security modes' of *Samba*, as set by the 'security' directive in the [global] section of `smb.conf`. Typically, you will see a line like this:

```
security = user
```



TutorialSamba

◀◀ which means that users will be authenticated against *Samba's* local account database, and access to shares will be controlled on the basis of that identity. There are three other possible security settings: share, server, and domain.

'SHARE' LEVEL SECURITY In this mode, which comes from the early days of Windows networking, users are not required to identify themselves but are simply required to supply a valid password on a per-share basis. Share level security is like having a key to a door. The door opens because you have the right key, not because it knows who you are. This mode sits very uncomfortably on top of the *Samba*/Linux security model because, as we've already seen, *Samba* needs to figure out a Linux UID to run with when it's accessing the files in the share. For share level security, it has a complicated and confusing set of rules for doing this. We won't go into the details, because share level security is pretty much obsolete.

'SERVER' SECURITY MODE Server security mode is similar to user mode, except that *Samba* will first try to authenticate the user by consulting another SMB server such as an NT server (or another *Samba* server). This is useful if you have several *Samba* servers on the network and want to consolidate your user account and password data onto a single machine. For example, the directives:

```
security = server
```

```
password server = SATURN
```

tell *Samba* to forward all authentication queries to the machine SATURN.

'DOMAIN' SECURITY MODE Finally, in domain security mode *Samba* will try to validate the username/password by passing it to a Windows NT Primary or Backup Domain Controller. Additional setup is required for this to work. In particular, a machine account (also known as a 'trust account') needs to be created for the *Samba* server on the domain controller. This is beyond the scope of our current discussion.

Note that whilst 'server' and 'domain' security modes delegate the authorisation process, they do not obviate the need for the client user to have a Linux user account.

Authorisation

The discussion so far has all been about authentication. What about authorisation? As we've already seen, authorisation checks are made at two levels, first by *Samba* itself, under control of

directives in *smb.conf*, and second, by the underlying Linux filesystem. Here, we'll focus on *Samba's* authorisation directives.

IP-BASED ACCESS CONTROL One of the simplest forms of access control is done not on the basis of user identity but on the basis of the identity (or more specifically the IP address) of the client machine. For example, a directive in the [global] section of *smb.conf* like this:

```
hosts deny = ALL EXCEPT 192.168.0. 127.
```

will disallow connections from all hosts except those on the 192.168.0 network (since this is one of those non-routable 'private' IP network addresses, this entry presumably refers to the local internal network). The "127" entry also allows connections from the local machine using the loopback address. This is important because although it's unlikely you'd want to connect to a *Samba* share on your own machine, some administrative tasks like changing *Samba* passwords, and running the configuration tool SWAT, require a connection to the local *Samba* server.

USER-BASED ACCESS CONTROL You can explicitly specify which users are allowed to access a share. For example, the share definition:

```
[payroll]
```

```
path = /accounts/payroll
```

```
read only = no
```

```
browseable = yes
```

```
valid users = david mary
```

allows access to the [payroll] share by the users david and mary. Alternatively you can use an entry such as

```
valid users = @accountants
```

which allows access to anyone in the Linux group accountants (or in the NIS netgroup of that name, if the system is using NIS). If there is no **valid users** directive, the default is to allow *all* users to connect. This isn't as open as it sounds, because in reality it's the underlying Linux file permissions which will limit who can do what with the files in the share.

GUEST ACCESS I made the point earlier about "no authentication without authentication". Well, there's an exception to this rule. *Samba* allows you to serve shares to users who are not able to authenticate – the so-called guest users. Examining how to set up guest access is an interesting exercise in understanding how *Samba's* authorisation mechanisms and the underlying Linux security model work together.

What we'll do here is to establish a guest share called 'blackboard' corresponding to the directory /home/bboard. Now, even for unauthenticated guests, *Samba* must have a Linux user identity to adopt when accessing the share. In this case we'll create a 'mythical' user called sam (Mythical in the sense that there is no warm pink body of this name; and no-one can actually log in to Linux as sam. The account exists simply to give *Samba* an identity to use when it is accessing the share on behalf of guests.)

First, we'll create the Linux user account like this:

```
# useradd -c "Samba guest" -d /home/bboard -s /sbin/nologin sam
```

Specifying a 'shell' of /sbin/nologin ensures that no-one can log in to this account; consequently, sam's password is essentially irrelevant. (Guest logins aside, it's worth pointing out that if you're creating a user account on Linux for the sole purpose of allowing access to shares through *Samba*, rather than (for example) allowing the user to log in and type commands at a shell, it's a good idea to disable the login as shown above.)

Samba resources

Useful links for more information

Samba is probably one of the best documented Open Source projects around. The man page for *smb.conf* is extremely detailed, and is also available as part of the SWAT package as an HTML file (*smb.conf.5.html*). There's also an extensive collection of HOWTO pages, collected together as a single PDF file. On a Red Hat system it's in /usr/share/doc/samba-xxx/docs, where xxx is the version number. All of this, and more, is also available online, at the main *Samba* web site

www.samba.org. Just follow the 'documentation' link. You can also download source and binaries of *Samba* itself from this site.

If you prefer something you can read in a hammock, I'd strongly recommend O'Reilly's *Using Samba* by Ts, Eckstein and Collier-Brown (ISBN 0-5960-0256-4). Not only is this one of the most authoritative books on *Samba*, it's also one of the most current, with coverage of *Samba* 2.2 and the upcoming version 3.0. Make sure you get

the second edition published February 2003 though – the original is getting a bit dated.

Finally, I was intending to provide a link to Daniel Robbins' excellent series of *Samba* articles that were on IBM's developerworks web site, but I just checked, and it seems they are there no more. This site does, however, offer a wealth of quality, technical tutorials and you might want to check it out. Go to the www.ibm.com/developerworks site and follow the 'Linux' link.

The `useradd` command will also create the `/home/bboard` directory and set it to be owned by `sam`. It looks like this:

```
# ls -ld /home/bboard
drwx----- 3 sam sam 4096 Jun 15 11:27 /home/bboard/
```

Now we need to add two lines to the **[global]** section of `smb.conf`, like this:

```
[global]
...
map to guest = Bad User
guest account = sam
...
```

The first line tells *Samba* that if a user tries to connect who does not have an account on the server, it should map the identity of the user to that of the guest account. The second line specifies the identity to use for the guest account – the mythical user ‘sam’ we just created. (As for all the directives, there is a compiled-in default for this; in this case it’s usually the user account ‘nobody’.)

Finally, we will be needing a section in `smb.conf` to define the share itself:

```
[blackboard]
path = /home/bboard
guest ok = yes
writeable = yes
```

The **guest ok** line means just what it says – it’s OK to allow guest users to connect to this share. The last line makes the share writeable – probably not something you’d really want to do for a guest share in the real world.

After making these changes, you’ll want to tell *Samba* to re-read the file. You can do this by sending a `SIGHUP` signal:

```
# killall -HUP smbd
```

or, on Red Hat, by using the handy **service** command:

```
# service smb reload
```

Now, if all is well, you should be able to connect to the share as a guest. To test this I created a user account called `sara` on a Windows 2000 client. `Sara` does not exist as an account on my *Samba* machine. I was able to connect to the share, and to read and write files there. Interestingly, if you look at the files that `sara` created, you’ll see that they’re actually owned by our mythical guest user, `sam`.

What happens if you connect to this share as a user that *does* have a valid account on the server (ie a user with entries in `/etc/passwd` and in `smbpasswd`)? To test this I created another user `mike` with accounts both on Windows 2000 and on the *Samba* server. Whilst `mike` is able to see the `blackboard` share, he won’t be able to list the files in it. Why is this? In this case, *Samba* is happy to allow `mike` to connect to the share, but the underlying **rw** permissions on the `/home/bboard` directory won’t allow him to access the file. In this case, *Samba* is now running as ‘mike’ not as ‘sam’.

Fancy footwork with Samba

This introductory tutorial has focussed on the use of *Samba* as a file server and describes functionality which has, for the most part, been around in *Samba* for quite a while (since version 2.0). Version 2.2 added and consolidated several other capabilities which allow *Samba* to take over a number of additional roles in Microsoft networks. Principally, *Samba* can act as a PDC (Primary Domain Controller). This includes the authentication of Windows 95, 98, NT, 2000 and XP users against its security database. Note, however, that Windows 95, 98 and XP home edition don’t

perform a full domain login in the way that Windows NT, 2000 and XP Professional clients do.

Samba also supports domain logon services including logon scripts, roaming profiles and system policies. Logon scripts are scripts (.BAT or .CMD files) that are executed on a client when a user logs on to a domain. A common use is to automatically map shares onto network drive letters. The scripts are held in the domain controller, but run on the client. Roaming profiles are somewhat inappropriately named – it’s the user that can roam; the profile stays put, that’s the whole point. Profiles define user preferences such as the desktop configuration and what’s on the menus. Roaming profiles allow users to log in on any client and download their own profile from the domain controller.

One thing that *Samba* won’t do is to act as a Backup domain controller (BDC) against an NT-based PDC (or *vice versa*). This is because it doesn’t support the protocol which Microsoft use to synchronise the security database between PDCs and BDCs. You have to remember that Microsoft don’t publish their protocols, nor do they make available the source code for their implementation. Writing code to interoperate with Windows on the network involves a good deal of packet-sniffing and reverse engineering. It’s a testament to the hard work of the *Samba* team that anything works at all!

The winbindd service


As we’ve seen, although *Samba* can delegate authorisation to an NT domain controller, it still requires all clients to have valid Linux accounts. These can be in the local `/etc/passwd` file, or be served from the NIS `passwd` map, if you’re using NIS. Although NIS allows you to centralise the user account database, you still need to create user accounts for each of your Windows client users, and with a large user population this is a chore system administrators would prefer to avoid. A fairly recent addition to *Samba*, the `winbindd` daemon, extends the range of mechanisms which Linux can use to perform user name lookups by allowing an NT server to be consulted.

This requires changes to the entries in the name service switch file so that the resolvers (the library functions which actually perform the user name lookups) know where to look. For example, an entry in `/etc/nsswitch.conf` like this:

```
passwd: files winbind
```

tells the resolvers to first look in the local `passwd` file and then to consult the `winbindd` daemon. For details, see the `winbindd` manual page.

On the horizon

At the time of writing this tutorial, the beta release of *Samba* 3.0 has just been announced. By the time you get to read this, a production release may be out. The authentication system has been almost completely re-written in this release. There is support for Active Directory (according to the release notes, “this release is able to join a ADS realm as a member server and authenticate users using LDAP/Kerberos”). However, *Samba* 3.0 cannot function as an Active Directory domain controller, and the restriction of not being able to synchronise between NT PDCs and BDCs remains. There is improved support for UNICODE (16-bit character sets). There’s a new ‘net’ command, designed to look like the net command in Windows/DOS, and a slew of other improvements. This release looks set to move *Samba* a good deal further down the road to provide a complete Linux/Windows integration solution. 

3D RENDERING

Realsoft 3D

Nick Veitch introduces you to this powerful 3D software, an exclusive version of which is on the *LXF* coverdiscs.



Realsoft 3D is a full-featured 3D modelling and raytracing package. With this software you can create, animate and render highly realistic or very stylised images. The special version that is included on our coverdisc has a few restrictions – mainly in the size of texture files you can use, and the size of images you can render, which is set at a maximum of 640x400 – plenty to get a feel for the application and create some great images and animations.



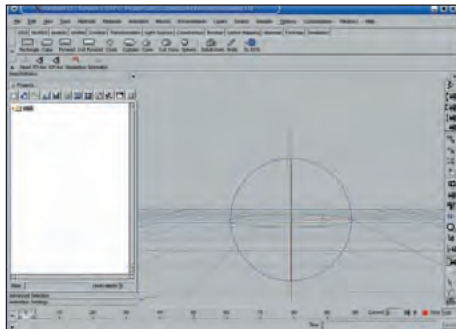
Installing

You'll find the *Realsoft 3D* directory on the coverdiscs in the Graphics section (on CDB if you read the CD version of the magazine). The file ending '.sh', is a script which unpacks and installs the *Realsoft 3D* software for you. You should be root to run this script if you want to install the software where anyone on your target system can use it. Once installed, you can run the system just by typing **realsoft3D**


Using NURBS

If you are a keen modeller, you'll maybe know that NURBS is an acronym that stands for NonUniform Rational B-Splines. To


everyone else, that just means they are surfaces constructed from curves. This short project will get you started.



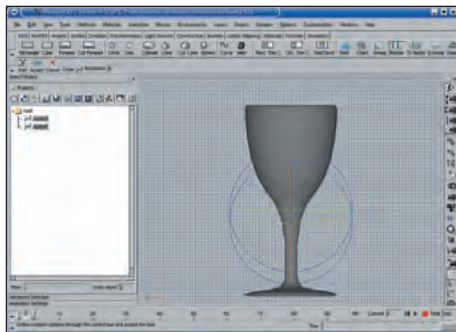
1 Probably the all time classic intro to using NURBS is revolving a curve into a wine glass. It's the 3D equivalent of a "hello world" program. So, being sticklers for tradition, that's what we'll start with. When you start up, the screen will look just like this. Click on the NURBS tab to get to the NURB related modelling tools.



2 Before we start though, we'll want to change a few settings. The default view is perspective, which is good for getting an overall view, but lousy for drawing accurate models like we are about to do. From the Windows menu, choose View Properties. In the camera panel, click on Orthographic, for a straight-on view. You can also move to the Input tab and activate Grid Draw, and perhaps even Grid Snap if your drawing is shaky. For this project, it's probably best to leave the Snap off.




3 Now we need to draw two curves. Select the curve tool and draw a straight line, top to bottom. Click for the first point, then hold down shift and click for the second point. (Shift will ensure you get a straight line). Right-click now and choose accept to accept the curve. This will be our rotational axis. Draw a second curve against this axis to represent a half cross-section of the glass. Click and drag to define the curves. It doesn't matter if you don't go right up to the axis, as this bit will be filled in.



4 With both curves created, select the Rotate tool from the menu bar at the top. Whenever you select a new tool, some directions appear in the status bar at the bottom of the screen which you might find helpful.

Now select (easiest to do this from the palette on the left of the window) first the axis and then the outline curve. As soon as you do this, you'll see a representation of the curved solid. At this point you can make changes in the lower of the two toolbars. Enter a different number in the 'Resolution' field to determine how many subsections are used for example. Click on accept to confirm the tool. At this point you might want to view the object in shaded mode – click on the icon that looks like a shaded diamond, three up from the bottom in the panel on the right.



5 A new object is now created. In the panel on the left, you'll now see an entry called 'mesh' with a number after it. This is the combined object. A powerful feature of *Realsoft 3D* is that the rotated object is still editable. Select the curve object and choose Edit from the tool menu. You can now adjust the profile curve!

Realsoft 3D Manual

We can't possibly explain everything there is to know about *Realsoft 3D* in this tutorial, but you'll be pleased to discover the software comes with a rather excellent manual, including further step-by-step instructions.

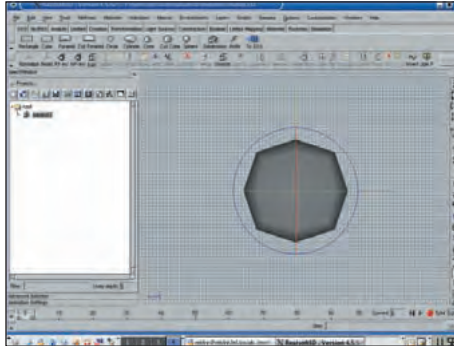
It's installed with the software. To access it, simply choose User Manual from the Help menu, and it will pop up in a browser. The buttons to the left and right in the coloured bars you can see in the screenshot will help you navigate through the manual to the relevant section.




SDS modelling

SDS stands for subdivision surfaces. These are another powerful way to create objects – you start off with a primitive shape, then subdivide and change it to create the object you want.


This time we'll look at modelling a different sort of glass in a completely different way, just to show there's always more than one way to achieve things!



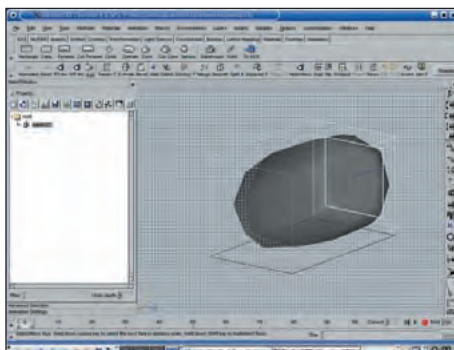
1 Click on the SDS tab to get the SDS tools visible. There are a number of shapes on offer, but we're going to use a cube. This may sound odd, but when you draw the cube, you'll see it doesn't actually look very cubic at all. Drag out a roughly square shape in the middle of the screen to get started.



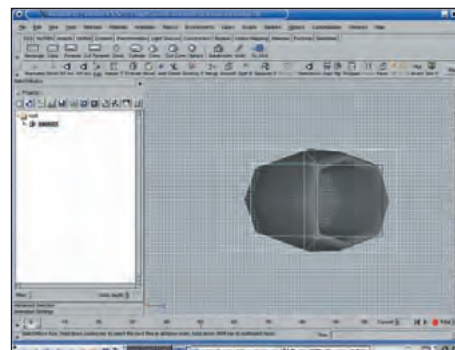
2 Now we need to edit the top face of our cube. Click and hold on the Edit button in the lower toolbar and select Faces from the drop down menu. This allows you to edit individual faces from the SDS object.




3 At this point it is probably worthwhile rotating the view a little so you can see what's going on. The right-hand palette contains lots of useful tools. The four camera icons underneath the eye can be used to rotate, zoom, pan and change perspective. Use these to get a view somewhere between face-on and side-on. If your object doesn't look like this, click on the Edit tool again and go back to object. Locate the point at the end of the blue bar protruding from the object and click and drag it to stretch the object out.



4 With Edit mode back to faces, choose the front face and select SubdivF from the toolbar. A blue bar with red endpoints now protrudes from the face. Click and drag it to alter the properties of the object. You want to create a rounded cylinder appearance for the top of our glass, so a bit of trial-and-error may be needed.



5 Again with the front face selected, click on the Extrude button in the toolbar. In this case we will be extruding back into the object. Grab the endpoint again and push it back down the cylinder. This hollows out the centre to create the inside of our glass.




6 To finish off the object, edit the bottom face and do a Subdiv F again to create a flat bottom (a glass is no good if it can't stand up!). To be able to usefully use this object in other scenes, it would be handy if it was the right way up! Choose Object Editing mode and rotate it (grab the coloured 'wheels' with the mouse and spin it). If you are happy with the object, you can smooth it out with the smooth tool in the menu. »

TutorialRealsoft3D

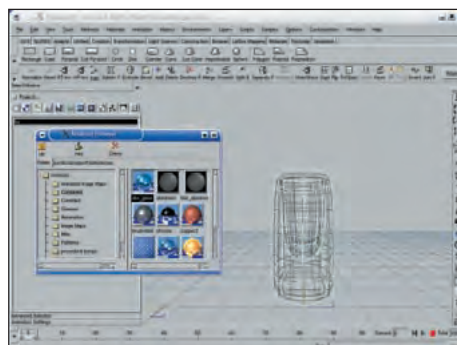
« Making a scene

We have objects, but what about making pretty pictures? There are a lot of things still to cover – materials, lighting, cameras and


render settings – but we're going to have a whistle-stop tour of the lot here. Remember to use the manual for more information.



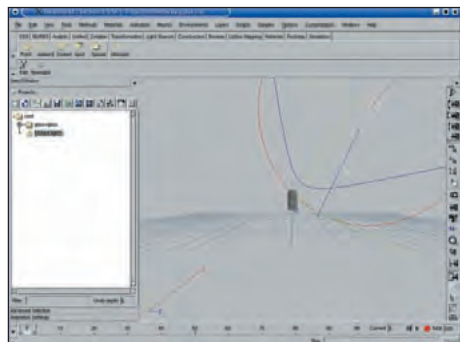
1 Load up your saved object and we can get cracking on building a scene. It helps if your object is the right way up and ready to go. If you didn't manage to create a decent object, or you want to try something else, you'll find some sample objects that you can start from in the `/usr/local/real3d/objects` directory.



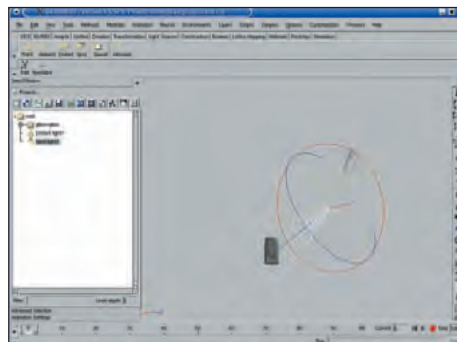
2 The first thing we need to do is add some material to our object. Select your object and click on the chequered square tab in the left-hand panel. This brings up the material list, which is usually pretty empty. Right-click in the panel and select 'Browse Material library' from the pop up menu. This brings up a categorised browser. Click on 'Combined' and you'll see some materials in the right-hand side of this panel. Just click and drag the material you fancy on to your object to assign it.



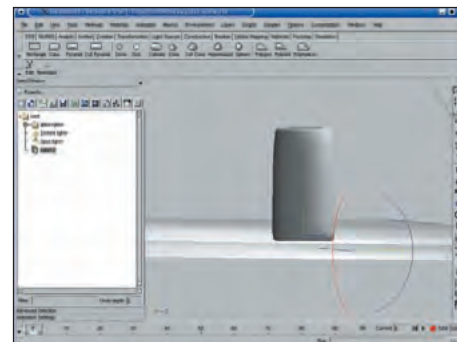
3 Click on the Ball tab in the left hand panel to get back to the object view. You should see your object has now been combined with the material you selected. There is obviously more to materials than this – you can assign multiple materials to the same object for example, or adjust any of the usual parameters for specularity etc. You'll find more details on this in the manual.



4 Now we need to add some lights to our scene. To create some general illumination, we'll first add a distant light. Click on the Light Sources tab in the toolbar and select distant. Zoom out and click to place the light, fairly high up and not quite overhead of the glass. Click once to place the light, then again in the general direction it is shining.



5 To pick out highlights on the glass, we'll also use a spotlight. This can be placed quite close in – the light source itself won't show up in our render. You may want to rotate the view a little to prevent the lights being on the same plane. Click once for position, once for direction and then to define the inner and outer cone of the light (the falloff is between these two areas).



6 Our scene is no good without a surface for the glass to stand on. For this simple scene, we'll create a floor or tabletop for it to rest on. Choose a suitable view and then create a cube from the Analytic tab. Drag out the cube to form a table.

Getting Realsoft 3D

The standard price for the full version 4.5 of *Realsoft 3D* for Linux is 300 Euros. Educational discounts are also available. To get more info on this product, visit the Realsoft website at www.realsoft.com

Error messages

There are still a few issues that you may notice when using the Linux version of *Realsoft 3D*. These won't prevent you using the software at all, but may result in some confusing error messages:

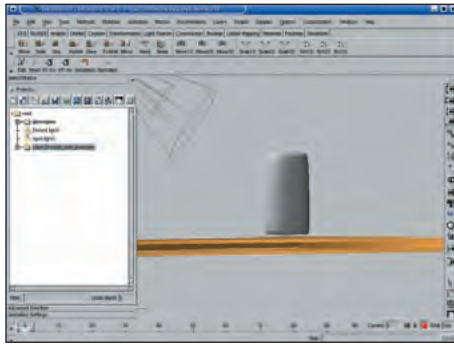
I GET GLIBC LIBRARY ERROR MESSAGES

The easiest way to solve this glitch is to run *realsoft3d* from the directory in which you installed it

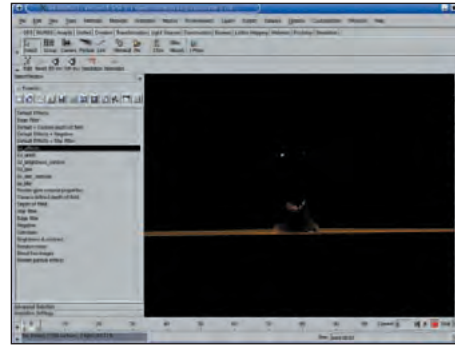
(`/usr/local/realsoft3d` by default). Change to this directory before running.

I GET ERRORS RELATING TO LIBAVIPLAY

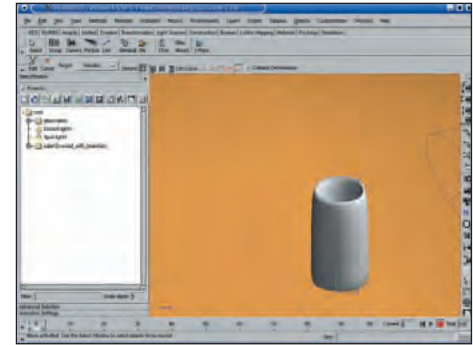
The internals for displaying avi animations are built against a different version of your player. It's safe to just ignore these messages, and get around the problem by using an external player for showing animations.




7 Switch to the op view and size the cube by dragging the handle to create a big enough table to fill the view. Assign a material as before – Wood is a good one to use (it can be found in the patterns folder).




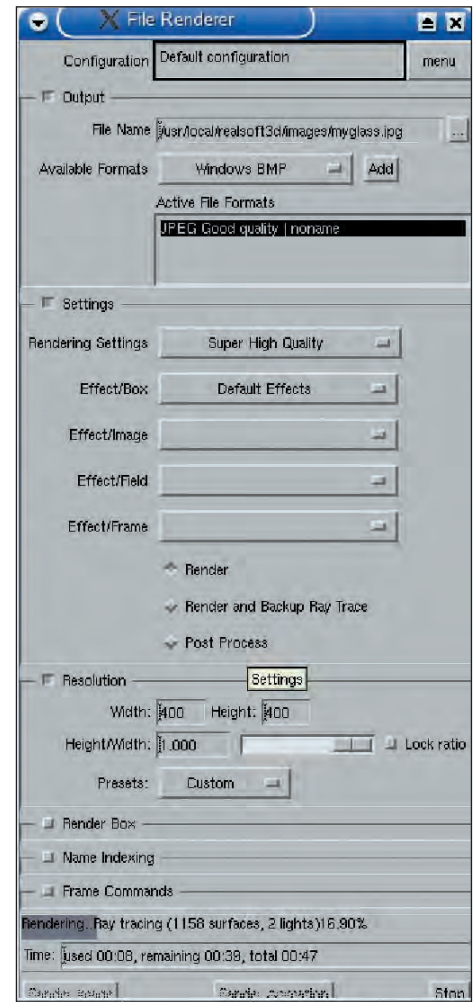
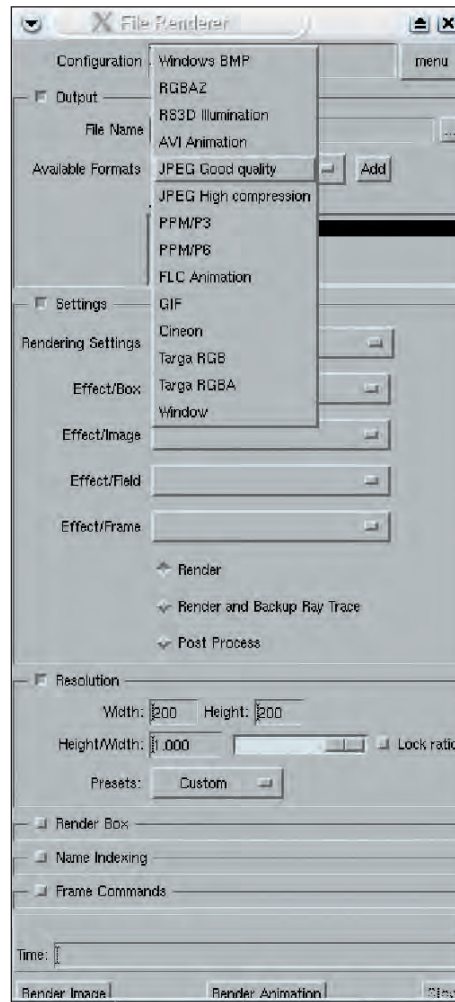
8 Now if you select the Render button (bottom of the icons on the right-hand side), you'll be treated to a quick preview of the scene after a minute or so. It doesn't look too good at the moment. The thing with glass is that it is transparent, so unless there is something for it to reflect or something behind it, it'll just look black, as you can see in the example above.



9 Move the view up and rotate it so you are looking down at the glass and table. To render the scene properly, we need a camera. Click on the Creation tab and click Camera. Without moving or clicking anywhere in the view, choose Accept from the lower toolbar to create the camera with exactly the same properties as the current view.



10 Do another test render to make sure the scene looks good. You might need to adjust the wood material to get a good texture underneath. You can also check the position of the lights and the shadows now.



The final rendered scene.

11 Choose Render from the File menu to bring up the rendering interface. The first thing we want to do is define the type of image to be created and where it is to be saved. *Realsoft 3D* defaults to a BMP image. Choose JPEG high quality from the available Formats box and select Add. Right-click on the BMP entry and select Remove. Change the name in the filename box to something sensible (and change the path if you like – open a requestor by clicking on the adjacent box).

12 The rest of the settings should be fine. You may want to turn up the Render settings to Super High Quality – this will produce the best glass effect. Without it, the inside of your glass may look a little dark. 200x200 is probably fine for this first render. Click on the render button and away you go. The status bar will let you know what's going on. **LXF**

ILLUSTRATION BY ROB DEBRICHY - www.blender3d.org

KEY FRAMES AND BONES

Get animated!

PART 5 Jono Bacon shows us the first steps towards our own computer-generated animated masterpiece *à la* Toy Story...

Last month we were looking at some of the texturing and materials facilities within *Blender*. Using these techniques, combined with techniques from previous parts of this series, we can create quite realistic-looking scenes with relative ease. Although we can use these methods in a variety of ways, all of our projects have been limited to still images. This month we are going to take *Blender* further and take a look at

some of the concepts of performing basic animation. By the end of this article you will be able to create simple animations using two main techniques.

Animation in *Blender* uses a variety of methods to move our objects in desired ways, but one of the main methods is the use of Key Frame animation. Key Frame animation involves creating a number of frames (key frames) in which the positioning and attributes of your scene are stored. When you set a number of these key frames in a sequence, *Blender* will calculate the animation between the frames. As an example, let us assume that our sequence has 60 frames. On frame 1 you may have a

cube in the centre of the scene. At frame 20 we can move the cube over to the left of the scene and at frame 60 we can move the cube over to the right of the scene. When the animation is played the cube will move from the centre of the scene to the left, over the space of 20 frames. The cube will then move from the left to the right of the scene over the space of 40 frames; the animation filling the gaps between the key frames.

Key frame animation is certainly the simplest method of animating your scene, but other techniques can be used to specifically determine how your objects move. We will be focusing on one of the more advanced methods of moving objects later in this article where we will add bones to our models.

Getting started with key framing

To get us started with key frame animation, we need to create a scene. Any scene will be suitable for this example, and I will be using a table scene that I built a while back. The table is in the corner of a room with a number of objects on it. Create your scene, appropriate lighting and materials and save it. My scene is shown in **Fig1**.

Key frame animation involves a few basic steps. First we must choose the frame that we want to make a key frame, next we move our scene appropriately and then we store our key frame. The first step is simple, and along the main button row (where the Light, Materials, Display and other buttons are) you will see a button with the number **1** on it. This is the frame button, and by clicking on the right side of the button will increase the number, and clicking on the left side will decrease the number. Starting with frame number **1** our scene will need to have each object stored in the key frame. It is best to do this on a per-object basis, and we will begin with the blue bottle in the scene. Select the bottle with the right mouse button and press the **I** key. Now select **LocRotSize**, and the location of the bottle as well as its current rotation and size will be stored in frame 1. Now move frame button to frame 20 and move some elements of the scene. The choice of what you move is entirely up to you, but pretty much any kind of change is acceptable. This includes moving, rotating, resizing and other operations.

In my scene I will move the blue bottle object upwards so it is hovering above the table. When the object has been moved, and while it is still selected, press the **I** key and select **LocRotSize** again to store the bottle's positioning in frame 20. We can now test our animation by moving the frame button back to frame 1 and by pressing **Alt-A** while the mouse cursor is hovering over the camera view. You can perform the animation preview while hovering the mouse cursor over the other views, but the result is not very interesting, so it is best to do it over the camera view. You can stop the animation by pressing **Escape**.

We will now move some of the other objects around on the table. We will keep the bottle in the air for frames 20 – 40, and move the sheet of metal under the bottle. First we will need to store the location of the metal in frame 20 where it currently is, and then move to frame 40 and store the location of both the bottle and the metal. Continue moving objects around your scene at different key frames, and include some rotation and scaling.

When you are happy with the animation in your preview (with **Alt-A**) you can render the scene. Switch to the Display buttons (**F10**) and click the **ANIM** button. When the scene has rendered, you can view the animation with the **PLAY** button. Within the Display buttons are a variety of other options that are worth looking at when configuring your animations. The first and most

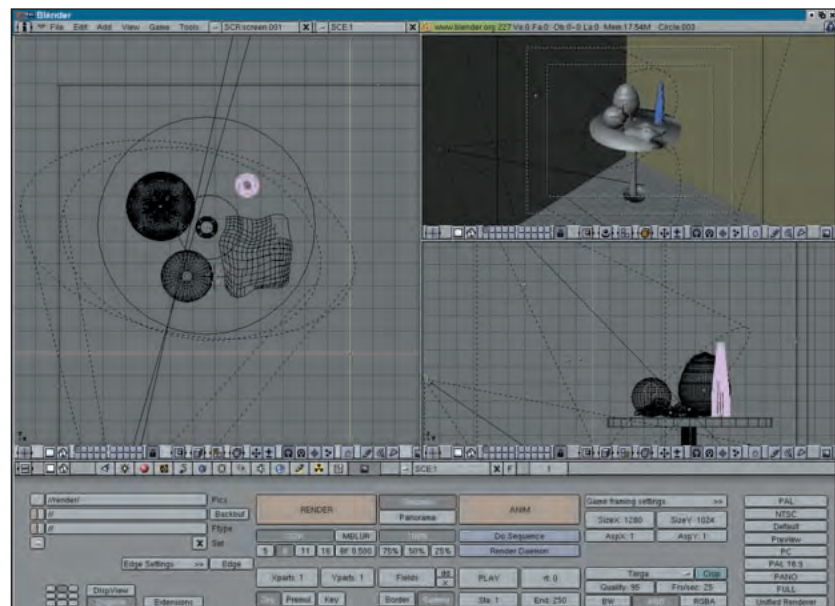


Fig1: A default scene that we will be animating.

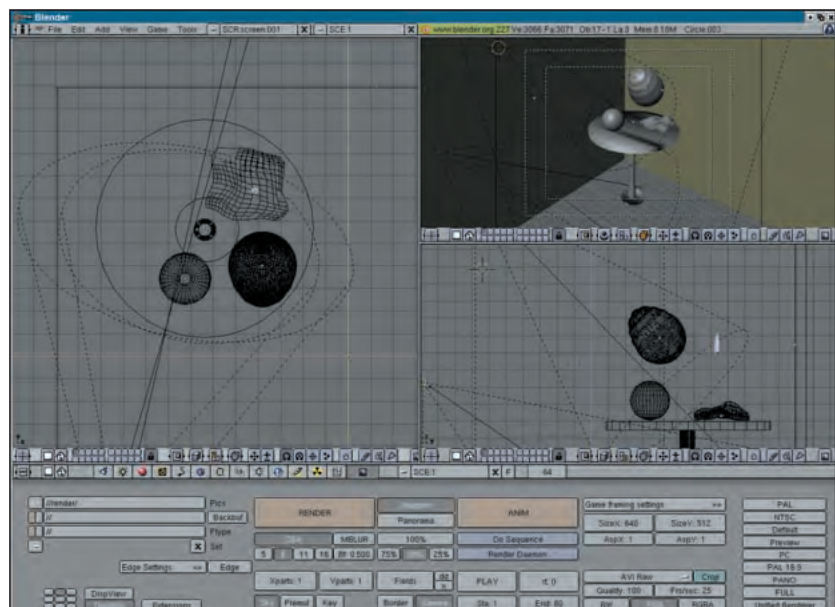


Fig2: Previewing a key frame animation.

important button is the drop down box with Targa written on it. This button lets you select what kind of file format you would like to save your animation as. The most popular is probably AVI RAW, which is a commonly supported AVI format, but produces fairly large files. Although large file sizes, AVI RAW will at least give you a source file to convert to other formats. Other options include the Start and End buttons which let you set the start and end frames and the Fr/Sec button which sets how many frames a second your animation runs at.

Bone-based animation

So far we have looked at frame animation and moving objects at the key frames. Although a good method of animation, the technique limits us in how much a particular object can animate. Apart from adjusting the size of the object, we have been fairly limited by using key frame animation. Bone-based animation takes things a step further and lets us literally give our objects



TutorialBlender

◀ bones (a skeletal structure). These bones allow us to move a bone and the relative section of the model will deform and bend with the bone. Using a combination of bones gives us the ability to have quite flexible animated sequences when combined with key framing discussed earlier.

To get us started with adding bones, we will need another test project. Create a new project, and using a combination of extrudes and scaling, create an object that is long and thin,

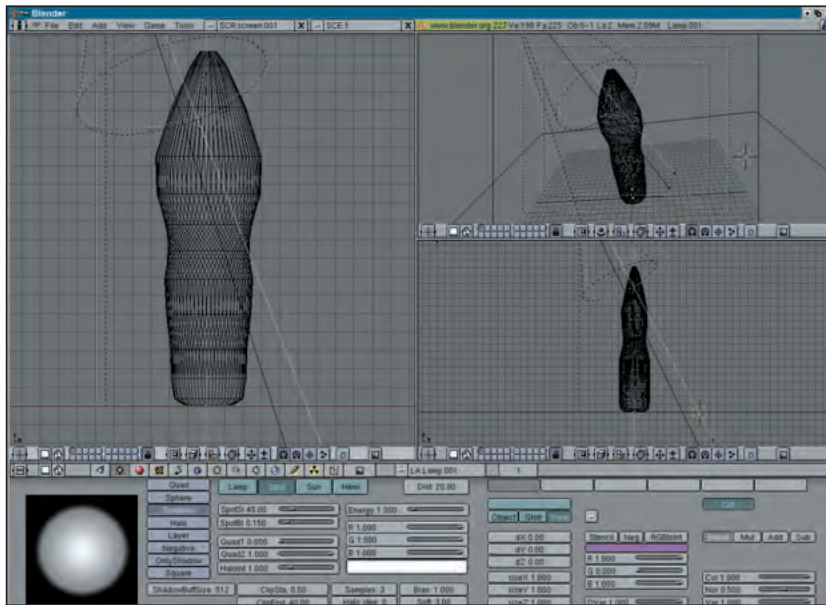


Fig3: The first stage of constructing a bendable object for adding armatures to.

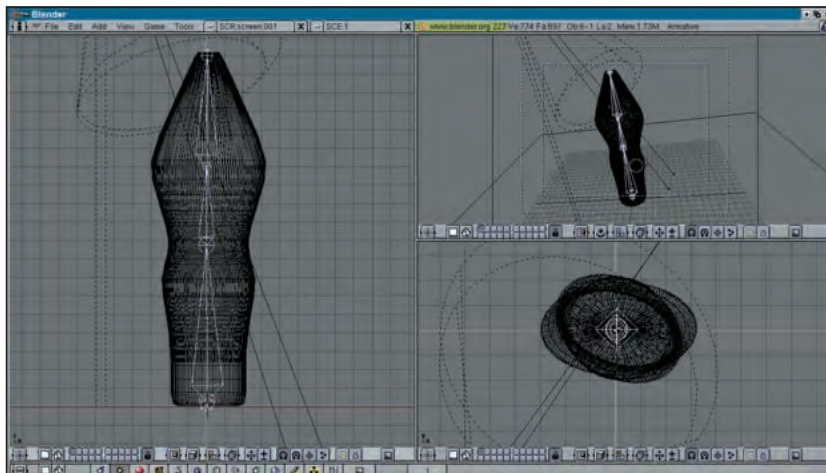


Fig4: Three armatures added to the object. These will act as the 'bones' of the object.

Yafray

Yet Another Free Raytracer

Although *Blender* has a built in rendering system, the recent releases of *Blender* are allowing you to use your own rendering system. One of the most impressive (and as such one of the most most popular) is *YaFray* (Yet Another Free Raytracer). *YaFray* renders models and animation to near photo-realistic quality, and many impressive examples of *YaFray* (like the one shown here on the right) are available on the Internet. You can find out more about *YaFray* at the official site www.yafray.org/



similar to a skyscraper in design. Feel free to make your object as lavish as possible, and add suitable lights and textures. The object I have created is shown in **Fig3** on the left.

With the object I have created, the largest view is the Front view. We will use this view to add the bones to the right parts of the model. Before we add the bones though, we need to ensure there are enough vertices to let the shape bend more smoothly. Generally the more vertices the better, but too many will overload the model. To create more vertices on the model, we will use the Subdivide command. First, select the shape and enter Edit Mode (**Tab**). When in edit mode, select all vertices with the **A** key so they are all yellow. Now press the **W** key and select Subdivide. Keep repeating this process until you have good number of subdivisions, but not so many that your model is a mass of yellow dots.

Now, de-select all the vertices with the **A** key so they are all pink, and put the 3D cursor so it is in the middle of the shape in the Top view, and at the bottom of the shape in Front view. Bones in *Blender* are called Armatures and are likened to mechanical armatures found in machines. To add an Armature, bring up the Toolbox with the **Spacebar** and select Armature. As you move the mouse you will see the armature gain in size, and when you are happy, place it in the right position with the left mouse button. A new armature will then be automatically connected. Keep adding armatures until you are finished, and press **Escape** to stop adding them. As you can see from **Fig4** I have added three armatures to my object.

Now the armatures have been added, we need to tell *Blender* which vertices each armature will affect when it is moved. Luckily for us, in *Blender* 2.27 a new feature was added to make this process much easier and simply move the vertices nearest to each armature. This generally makes sense and is suitable for this example. To do this, we need to parent our two objects (the model and the armature chain) together. First select the model and then using **Shift**, select the armatures also. Now press the **P** key and select Use Armature. Now select Create From Closest Bone to finish creating our bendable model.

To actually move our object, we need to enter a mode called Pose Mode. In this mode we can move the armatures and the relevant parts of the model will bend and deform as required. This mode allows to begin constructing some animation with the model. To enter the mode, first select the armatures and ensure that only they are selected. Next press **Ctrl-Tab** and the armatures will turn Blue; this means you are in Pose Mode. In this mode you can then select an armature segment by right-clicking it, and it will turn a darker shade of blue. This is now the selected segment in Pose Mode, and you can use the rotation and scale keys to move the armature segment. Ensure you have a camera looking at the object, and you will see that the object will move.

Combining key frames & armatures

Now we have covered both techniques, we can combine them to create an animation using the armatures we have just made. First of all, select frame 1 with the frame button, and use the **I** key to store the LocRotSize of the armatures. As with our previous examples, what you want to do specifically with the animation is completely up to you, but we will cover a few sample ideas here. We will first sway the object from left to right. Move to frame 10 and then rotate the top armature left in Front, and then insert the key frame. Next, move to frame 20 and move the armature all the way to right, before inserting the key frame again. Using the same techniques, we can now move the middle armature just a bit before

Animation and the camera

Fly-by or walkthrough effects

In this article we have only looked at how we can animate particular objects moving. Objects are certainly not the only entity that can be animated, and one of the most powerful items to animate is the camera. Animating the camera is not just a case of moving the camera, but it can offer some real depth to your animations, and some impressive fly-by or walk-through effects can be obtained by animating the camera through a variety of objects in a scene.

To animate the camera in your scenes, you need to make use of key frame animation as we have used with

normal objects. Select the camera and insert a key frame for the camera, and then move the camera to the desired location on the desired frame. Insert another key frame and the camera will move over to that frame. It is important when dealing with animated cameras to be aware of the direction of the lens. If the lens stays in the same position but moves its location, the camera will pan over to the new location. If the lens moves but the camera stays in the same place, the camera will simply rotate around.

Dealing with camera animation also brings up a very

important issue – always be careful that your camera does not run through any of your objects (unless you want it to of course). Remember that in the *Blender* 3D world, objects can run into each other, and as such the realism of solid objects can be shattered unless you ensure that objects do not run into each other. The other major issue to remember is to sufficiently light your scene. When you move the camera, you may have only lit the scene for that single camera view, but when you start animating the camera, you will need the entire scene to be correctly lit and textured.

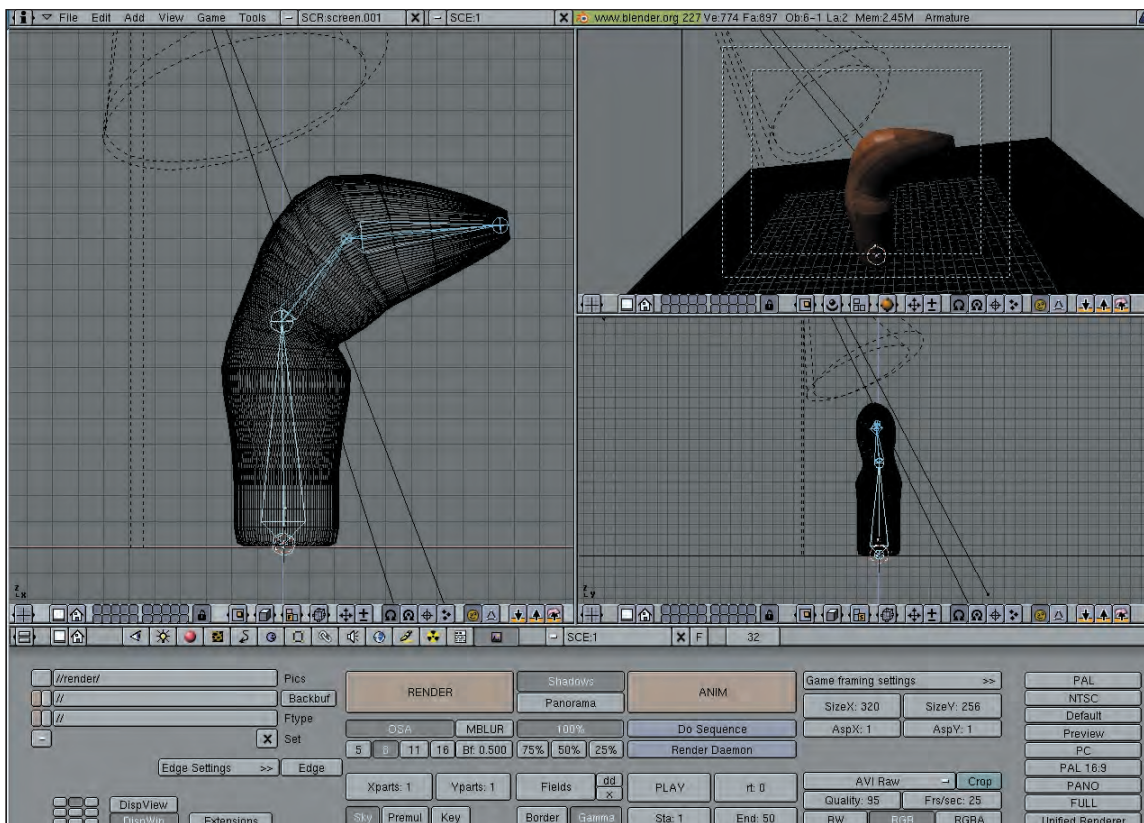



Fig5: Experiment with deforming your model to produce a variety of different effects.

the top armature, and the object will perform a shake, similar to a belly dancer. **Fig5** shows the model at one extreme of its deformation. Various other techniques can be tried such as swaying

the top of the object around in a circle-like fashion – this can be achieved by setting each quarter of the rotation as a key frame.

Conclusion

This issue we have taken tentative steps to scratch the surface of the complex subject of animation. Although we have covered some of the basic concepts of animation such as key framing, these techniques have a lot of flexibility in their combined effect. The technique of adding bones to a character, as an example, can be combined with key framing and an audio track to perform some kind of dancing animation. When you also combine into the scene elements that place the dancing object in a room, with volumetric lighting, realistic textures and other objects that move, we can already create quite convincing animated scenes. As you continue to learn about *Blender* and the wealth of functionality within it, you will no doubt make use of animation graphs, motion paths and other techniques. These techniques can be combined with the concepts discussed this month to help bring more life to your animations. 

Inspiration

When doing any kind of creative work, and *Blender* work is certainly included, it is often useful to have some creative inspiration to see what you can do with the technology and your imagination. *Blender* has recently been the platform that has created a number of 'homebrew' animated shorts that have been released on the Internet, and these shorts really show just what can be done with *Blender*. It is recommended that you take a look at these animations, and not only enjoy them for what they are, but also try to see how the creator has constructed the objects, lighting and other elements in the animations. This will help you when creating your own animations. The following are recommended:

Cyan Sun – <http://homepages.picknowl.com.au/leewj/cyansunweb/cyansun.htm>

Mindfields – www.artificial3d.com/mindfields/

NEXT MONTH

We will be taking a look at some of the more advanced topics in *Blender*, paying particular attention to more advanced modelling techniques and particle systems. We have now covered the major techniques in brief, and over the next few months we will look at some of the additional features that can be used to add polish to your scenes.

TutorialGIMP



IMAGE EDITING

Integrating text

PART 7 An image may be worth a thousand words, but often they still require the addition of text. You'll need to be focused if you expect to be able to edit GIMP text after you've created it, says **Michael J Hammel**.

Let's face it – Linux has not been, until relatively recently, well known for its font support. Scaled and blocky bitmapped fonts left users with sore eyes and archaic designs. Even so, Linux's problems only partly affected text management under *The GIMP*. The app has supported antialiased fonts for quite some time, for example. Yet the general lack of support under Linux for TrueType fonts has left *GIMP* in the lurch as well – both Linux and *GIMP* rely on the underlying X Window System to provide that support. In recent times the release of FreeType and updates to the X Font Server have added a world of fonts to both Linux and *GIMP* and the next release of *GIMP* promises to do even more. Until then, however, we need to make the tools available do the most they can.

GIMP 1.2 provides multiple tools for adding text to images. The default Text Tool, found in the *GIMP* Toolbox with the letter **T** as

its icon, provides basic font selection features. Unfortunately its font filtering is based on the fairly cryptic XLFD (X Logical Font Description) standards – known only to the most addicted Unix developers. A better font interface is provided by Dynamic Text, an alternative tool available as an option in the Text Tool's Tool Options dialog (double click on the Text Tool icon to open this dialog) as well as from the Canvas menu option Filters > Render > Dynamic Text.


A third option for generating text is the new FreeType plugin. Though not currently distributed in any stock *GIMP* distribution, it is available from the *GIMP* Registry (<http://registry.gimp.org/index.jsp>). The next release of *GIMP* (due out soon, we hope) will add FreeType as a standard tool and take Text management to places users have been dreaming of for years. For now, once you have FreeType installed, look for it under the Filters > Render > FreeType menu option.

In these tutorials we'll look at the services provided by the various Text management tools in the current version of *The GIMP*, version 1.2.5, along with the multitude of text plugins that extend those services.


Text Tool vs Dynamic Text

Long-lived applications often have features with little usefulness to current users. This is the case with the default Text Tool in *The GIMP*. While it served its purpose in the early years, it has long since been superseded by the Dynamic Text tool.

1 When you choose the Text Tool from the ToolBox and click in a Canvas window, the default Text Tool dialog will be displayed. This dialog has three pages, as shown in this image. The first page allows selection of a font based on a common name and offers a space to enter a single line of text. When you click on the 'OK' button, the text is rendered in the Canvas window as a new layer. The drawbacks to this default dialog are that you can only enter a single line of text and, once applied, you can't edit that text in the new layer. To make changes to the text, you have to use the Text Tool again and render text to a new layer, deleting the old one manually. On the other hand, if you have thousands of fonts installed, from many different font foundries, you can search for a particular font much more quickly using this dialog than either the Dynamic Text or FreeType dialogs. Only text scaling is offered here – no other text characteristics, such as shear, character spacing or rotation, are available. If you need these features you need to look into the Dynamic Text or FreeType tools.



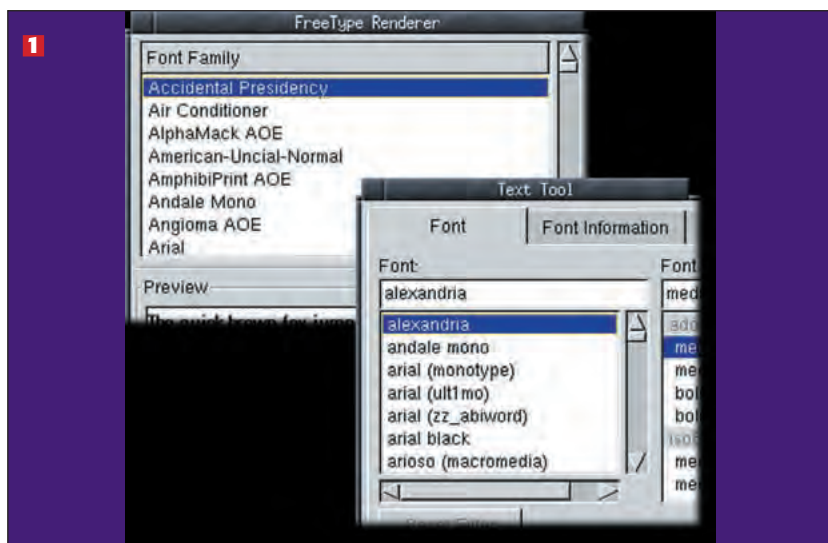
2 The Dynamic Text dialog is much more artist-friendly. To get this dialog you must first click on the Dynamic Text button in the Text Tool's Tool Options dialog, then click on your image to get the Dynamic Text dialog. The text input area (the white box at the bottom of the dialog) allows for multiple lines of text. Font selection here is by common name only which means there is one long list of font names. That can be cumbersome if you have a large number of fonts installed. Other features of this version of the Text Tool include text alignment for multi-line text, antialiasing of text when rendered into a new layer, text color setting, and rotation, line spacing and new layer alignment. The rotation option is in degrees but is not shown in the preview. Most important of all, this is the only text tool in *GIMP* that allows you to edit your text *after* you've rendered it to a new layer. Just make sure the text layer is active in the Layers dialog and then click on the image to open the Dynamic Text dialog. Changes to the text are applied when you click the Apply or OK buttons.



FreeType

Type 1 fonts have traditionally been the only fonts available to Linux users. Recently, TrueType support became available and, if you're distribution of Linux includes that support, *GIMP* now can make use of fonts in that format. The default and Dynamic Text tools, however, can only render (ie draw) those fonts – they can't manipulate the format itself prior to rendering it. That means they don't provide things like shearing, character spacing, hinting (detailed information for rendering of font characters) or kerning (character-to-character spacing). To get access to these features you need to grab the FreeType plugin from the *GIMP* Registry.


1 FreeType is available from the Filters > Render > FreeType menu option. Once the dialog is open, it searches for fonts it can manipulate, bypassing the normal font mechanisms used by the default and Dynamic Text tools. This means you may find different fonts listed by this text tool than you did in the other two. The image shows the first few fonts listed in both the FreeType dialog (upper left) and the default Text Tool dialog (lower right).



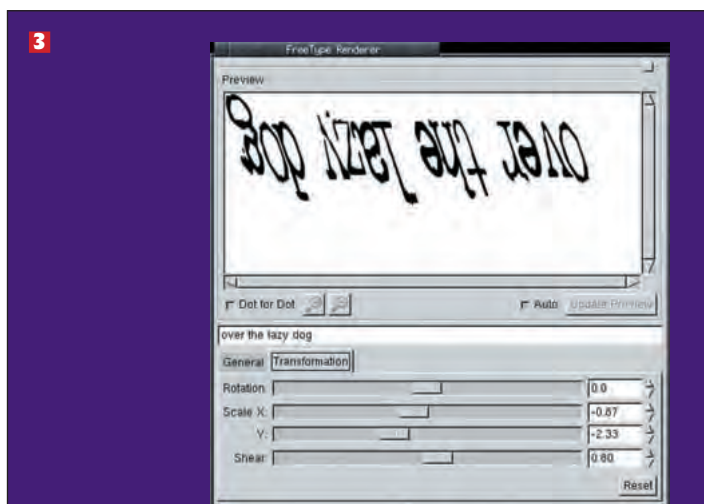
TutorialGIMP

FreeType (continued)

2 Options in the FreeType dialog allow you to quickly change the size and spacing of characters, even (in some cases) reversing the direction of the text. Here, the text “over the lazy dog” is shown in reverse. Kerning, hinting and antialiasing can all be enabled or disabled from here, with changes displayed in the preview immediately. Once rendered into a new layer however, the text cannot be edited. The new layer will always be placed in the upper left corner of the Canvas even if that area is not visible in the window (zoom out to see the new layer if you need to).



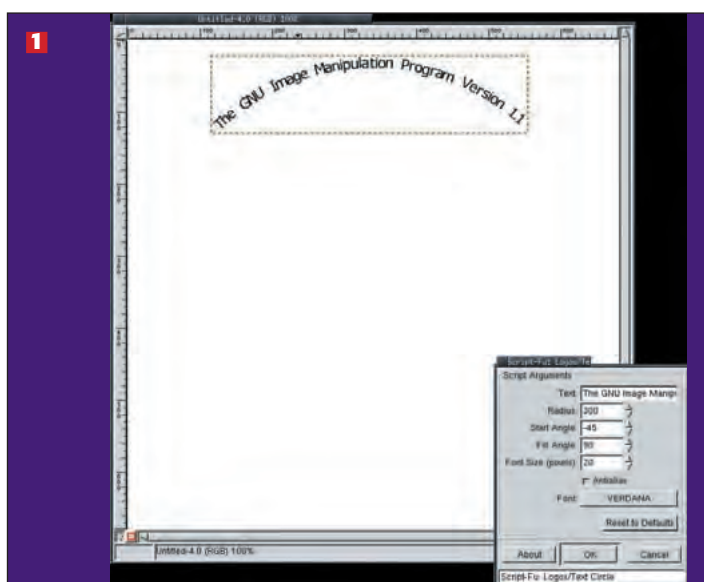
3 Text can be manipulated in multiple ways. In this example the text is reversed and flipped by scaling the X and Y directions to negative values. A little shear is also applied. All of this is done interactively so you don't have to wait for a new layer to be rendered and the content of your text can be changed. Unfortunately, with all this, there is no way to create multi-line text except to create each line as a new layer and position it manually.



Text along a curve

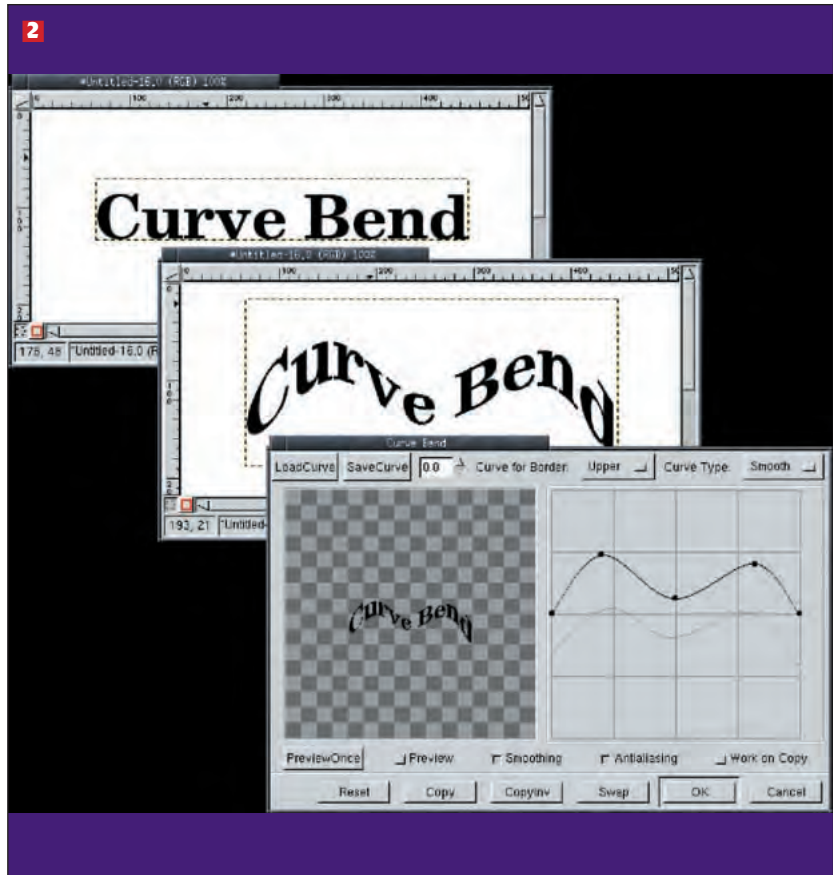
The three primary tools for managing text in *The GIMP* provide few text layout features. In order to manipulate the text layout you have to edit the text after its been rendered into a new layer. This is not the optimal solution since the pixels in the text are being changed instead of changing the position of the text, but with *The GIMP* 1.2 its the only choice. *The GIMP* 1.3 will offer features that allow real text layout to become possible.

1 The only real text-along-a-curve tool available actually shows up as a Logo script. The Text Circle script (Xtns > Script Fu > Logos > Text Circle) places the specified text along an arc of a circle with the specified radius. The Start Angle (in angles, counterclockwise) determines where the text will start and the Fill Angle determines how far along the arc the text should flow. You must specify a font to use since, by default, no font is selected. Some fonts won't work with this plugin, and since the letters are rendered one at a time and then rotated and translated to their appropriate position, the result may appear like smeared, overlapping lettering if your radius and angles are not large enough. In this example, the new image created by the Text Circle plugin is significantly larger than the actual text because the radius determines the size of the new canvas.



Text along a curve (continued)

2 An alternative to placing text along a curve is to shape text with curved boundaries. This can be accomplished with the Curve Bend plugin (Filters > Distorts > CurveBend). Use the upper and lower options to force the top and bottom edges of the text to follow curves. The result is that text appears to follow the curve. Note that using the upper and lower border options will cause your text to be fit between the two curves. This can cause text to be stretched or squashed. The end result is that you can lose some of the smoothness of the originally rendered text – a stairstep affect may appear, depending on the font used, after you've applied the curve. The trick here is to either use small curves, use a thin or rounded font or add some of the background colour to the text layer. Don't merge your background with the text or Curve Bend will apply the curve to the background as well and you'll end up with some transparent portions in your image. If you do that, add a new layer filled with the same colour as the background and drop it below your curved text. That will fill in the transparent areas. Finally, when working on text with Curve Bend, always leave the antialias and smoothing options enabled.



Logos

The GIMP and the Web both grew up at about the same time, so it's no surprise GIMP has been widely adopted for Web graphics design. One of the most common tasks for GIMP has been creating Web logos. Because it is so well suited to this sort of work, many scripts have been written to automate logo generation. Many of these scripts are included in the core GIMP distribution under the Xtns > Script Fu > Logos menu option in the Toolbox. More logo scripts can be found at the GIMP Registry or by searching the web.

1 The samples shown were created mostly from the default settings. All of the Script-FU logos have the same basic dialog layout: provide the text for the logo, select a font, select foreground and (at times) other colours or patterns, select shadow settings, and so forth. Most logo scripts will not have a default font chosen, so you need to be sure to at least choose a font, though most other defaults will normally suffice for experimentation. Some logo scripts will create a single image with no layers in it. Others will leave the layers in place so you can modify the logo manually after it is created. Often, dramatic changes can be made to the logo simply by changing the Layer Blend Mode for various layers in the generated logos image. Since these logos are written in the Script-FU language (which is a subset of the Scheme scripting language) you can copy the scripts from their home (usually /usr/share/gimp/1.2/scripts) into your \$HOME/.gimp-1.2/scripts directory and edit them to create your own custom logos. Unfortunately, Script-FU is not the easiest scripting language to learn. Logos can also be written in Perl if you have the Perl extension included with your GIMP package. [LXF](#)



FASTER CODE

Practical PHP programming

Optimisation is the blackest of black arts in programming, but when you don't want to go to the fuss of writing your own C extension, how else can you make your PHP code run faster? **Paul Hudson** overclocks his PHP scripts...

For the past three issues we've been covering the devilishly hard topic of writing your own PHP extension. Yes, it can yield huge performance returns for your site, but on the other hand they're hard to make, fussy to maintain, and, after all, this is a tutorial about writing PHP, right?

Performance, particularly on busy sites, can be critical – if you can speed up your code by 10%, that decreases your hardware load by 10%, saving you the need to upgrade. There are a number of ways you can improve the performance of your

scripts, and we'll be covering as many as have space for. We'll also be dispelling various myths about optimisation, and hopefully by the end of this mini-series you'll be confidently about to re-write fundamentally flawed algorithms, tune implementations of good algorithms, make your MySQL queries fly, and more.

Before we begin, I'd like to make it quite clear that optimisation is the process of *improving* performance of your code, whether to use less space or run faster – it's usually a trade-off. Optimised code isn't necessarily "perfect code", it's just better than unoptimised code.

Furthermore, there is rarely if ever such a thing as "too fast". In my spare time, I've been working on my own pet project: a gigantic online PHP strategy game. Actions occur at round end, which is triggered by a cron job every ten minutes. With a thousand dummy users in the system, round end takes over seven seconds, during which database writes are locked so that people can't make changes. In the past I've spent hours and hours just to cut one part of that round end from 0.002 seconds to 0.001 seconds per player – it might not sound like a lot, but as far as I'm concerned every single second counts, and that

Single quotes are faster!

Have you picked up the received 'wisdom' that using single quotes for your strings rather than double quotes can yield a substantial speed boost? If you have, I'm sorry to disappoint – this is one of the most commonly repeated performance myths about PHP. Whether you use double quotes or single quotes is pretty

much beside the point – yes, there is a performance boost to use single quotes, but it's much less than 0.01%, and it's generally just not worth the extra hassle. Many people use double quotes for everything, and that's fine – use whatever you feel most comfortable with, because it won't affect the speed of your script.

trimmed one whole second off round end. If you've tried out half of the recommendations here and find you've reduced the run-time for a script from four seconds down to one second, don't stop there – go for the fastest code you can get.

The easiest optimisation

There is one simple way you can double the speed of your server, and that is to install IonCube's *PHP Accelerator (PHPA)* – it's fast, free, and works perfectly on Linux. If you don't mind paying money and want another 20% extra speed boost, Zend's *Performance Suite* is the way to go – it's better-supported than *PHPA*, works with newer versions better, and has the edge on speed to boot. If you have a small/non-existent budget, *PHPA* is your best choice. If you've got a little money and really do want to push your PHP code to the max, *Zend Performance Suite* will be a purchase you won't regret.

If you want to read more about PHP accelerators, *LXF34* has a roundup. Since that article, *PHPA* hasn't changed versions, but *Zend Accelerator* got a big upgrade and became *Zend Performance Suite* – it now includes various other enhancements and speed ups that should increase the lead of *Zend Performance Suite* by a substantial margin.

Recently the *Alternative PHP Cache (APC)*, the cache that performed most poorly in our tests back in *LXF34*, has been transferred into PEAR, the PHP code repository – it's now being worked on by PHP developers. As such, it's now much better than it previously was, almost snatching the number1 spot. Well worth a try!

If you'd rather not install a full-blown PHP accelerator, consider at least using the *Zend Optimizer* – it's free, easy to install, and provides a handy speed-up for larger scripts.

Very basic optimisation

There are a lot of simple tweaks you can make to your PHP coding style that will produce small benefits, so we'll start there.

Firstly, take advantage of the fact that PHP allows you to post-increment (**\$i++**) and pre-increment (**++\$i**). The meaning is the same as long as you're not writing anything like **\$j = \$i++**, however pre-incrementing is almost 10% faster, which means that you should switch from post- to pre-incrementing when you have the opportunity, especially in tight loops.

Secondly, run tests to see whether your scripts can benefit from a switch to using references – very often you'll find that assigning by reference and returning by reference can speed up the handling of complex objects by quite a margin, as well as also lowering your memory usage. Note that from PHP 5, objects are reference-copied by default – I think that says a great deal about what the PHP developers think about referencing complex objects! Although this might not make sense if you're fairly new to PHP, switching to references across the board is rarely a good thing, which is why I recommend you run tests on your scripts before you make the change.

Thirdly, always set your PHP error level to the most verbose level, **E_ALL**. All too often people don't realise that PHP is outputting various complaints about variables not being set, etc, which you can just do away with entirely by cleaning up your code. While you're editing your `php.ini` file, it would also help to disable all the extensions you don't use – they're just chewing up memory otherwise.

Fourthly, upgrade to the latest PHP version. Fixes and tweaks to slow code are being made all the time, so it's worth taking

Take advantage of new features

Get updated!

Lots of functionality was added in PHP 4.3, and yet few people seem to be taking advantage of it. How would you, for example, read the contents of a file into a string? Like this?

```
$contents = implode("\n", file("test.txt"))
```

If so, you're using code that was considered very fast before 4.3 was released. However, in PHP 4.3 a new function was introduced called **file_get_contents()**, that takes a filename as its first parameter and returns as a string the

entire contents of the file. It's very fast as it's very optimised, and neatly solves a problem that people have had to implement for themselves.

If you haven't done so recently, take a good look through the PHP manual to see what's changed – lots of functions were added in 4.3, and lots more are being added in 5.0, including a **file_put_contents()** function that will write a string to a filename.



advantage of all the hard work going on. Note that as the fix to many PHP bugs is "upgrade to the latest version", this is generally always the best option – if you're running anything older than 4.3, upgrade; if you're on anything before 4.2, upgrade right now; if you're on anything before 4.1, what's taking you so long? Use the **phpinfo()** function to get your version number and information about activated modules.

Fifthly, don't repeatedly access the same element in an array unless you absolutely have to. Try running the following code on your own machine:

```
<?php
$START = time();
$foo['bar'] = "test";
for ($i = 0; $i < 10000000; ++$i) {
    if ($foo['bar'] == "test") {
        $j = 0;
    }
}
$END = time() - $START;
echo "Array took $END seconds\n";
$START = time();
$testvar = $foo['bar'];
for ($i = 0; $i < 10000000; ++$i) {
    if ($testvar == "test") {
        $j = 0;
    }
}
```

The Zend Performance Suite beats its competitors hands down, but will set you back a few clams.


```

<< }
SEND = time() - $START;
echo "Var took $SEND seconds\n";
?>

```

Running that took 38 seconds for the repeated array access, and 32 seconds for using a variable – the reason for this is that accessing array elements time and time again requires PHP to find the element inside the array, which is comparatively slow.

That's the easy stuff out of the way – you should be able to do all the above with no bother, boosting the speed of your scripts by a little at no cost. On to more complicated matters...

Compile right

One of the biggest advantages to using Linux is that we get to compile our software ourselves, and it really does make a big difference to the speed of your software. If you're able to, I suggest you compile *Apache*, *PHP*, and *MySQL* yourself using *GCC 3.3*, and do this with as many optimisations turned on that you have time to wait for.

Compress your output

HTML is a very wordy format – there's a lot of duplication in the form of HTML tags, but also in the main body of text. Furthermore, by default PHP will send text to *Apache* as soon as it's ready, which results in less efficient transfer of data.

The solution is to enable output buffering, and to use *gzip* compression for the buffers. Output buffering, if you were unaware, makes PHP store up its data into one big chunk, then send it to *Apache* all at once. Because all the data is kept together and sent all at once, PHP is able to compress it using *gzip* compression, which will generally reduce the content to around 33% of its original size (that is, a 1MB file will become roughly 300KB).

Not all clients support receiving compressed content (although nearly every browser made in the last five years will), and to handle that PHP will only compress data if the client can support it – this means you can enable compression, and not have to worry about old clients because PHP won't send them compressed data.

To enable output buffering for all your scripts, open up your `php.ini` file and set **output_buffering** to **1** and **output_handler** to **ob_gzhandler**. You'll find those keys already set in your `php.ini`, so just change the existing values. You should then make sure that you check your `phpinfo()` output to make sure output buffering is enabled correctly.

Don't use CGI

If you have the choice of using PHP as an *Apache* module or as a CGI, plump for the *Apache* module each and every time. PHP running as a module is about five times faster than PHP running as a CGI, because it no longer has to perform lots of startup code

whenever a script is executed – the PHP parser and modules are loaded when *Apache* is started, and kept in memory waiting for more scripts.

One other important reason to use the *Apache* module in place of the CGI executable is that PHP add-ons like *Zend Performance Suite* and *PHPA* don't work with either CGI or the CLI SAPIs, because they cannot cache across processes.

Debug your code

One problem with PHP is that, by default, if it encounters non-fatal errors messages it will just output them along with the rest of its output, which means that very often you don't notice the errors. While this might not seem like such a bad thing – after all the errors are non-fatal, right?

In the world of programming, one rule is fairly constant: code will run quickly until it has to handle errors. That is, errors in your code are likely to chew up five, ten, or even twenty times the resources that they should – you'll have noticed that *Apache*/PHP tends to shoot to the top of "top" and/or thrash your hard drive whenever it encounters a series of script problems.

You should thoroughly check the output your pages produce in order to make sure PHP isn't emitting errors behind your back. Alternatively, make sure error logging is turned on in your `php.ini` file, and check it regularly.

Cache your pages

PHP speed-boosters like *Zend Performance Suite* and *PHP Accelerator* work by caching your PHP code before it is executed, which provides a substantial speed boost. But what do you do if the PHP processing itself is taking longer than is acceptable? The answer is usually to cache the output of your pages after they have been output as HTML, which has two side effects.

Firstly, it will mean your pages will be static most of the time. This is obviously not an option for sites that need dynamic content all the time, such as shopping baskets and the like, but it is definitely possible for news sites where the content only changes once every few minutes. Even on messageboards, it's possible to print a message such as "your post will be displayed within the next five minutes", like <http://slashdot.org> does.

Secondly, it will mean you get the biggest possible speed boost, because you're only outputting HTML for the majority of the time. All PHP has to do is run a check along the lines of "if content is more than five minutes old, regenerate it, otherwise send the old stuff".

To implement this system, you need to modify your PHP scripts so that they check for the cached copy being out of date, and generate a new version if needed. Here's an example of how to implement this, with a ten-second cache life to make the caching fairly apparent – you'll probably want a five- or ten-minute cache life.

```

<?php
$cachefile = basename($_SERVER['PHP_SELF'], '.php') .
'.cache';
clearstatcache();
if (file_exists($cachefile) && filemtime($cachefile) > time() -
10) { // good to serve!
include($cachefile);
exit;
}

```

PHP 5 Beta 1

Available to download

The first PHP 5 beta is now available for download, and has many, many large changes from PHP 4. The final version is a long way off yet, but you should consider giving the beta a try out – it helps find bugs, and it also gives you an idea how well your code will work in the new *Zend Engine*.

Please note that the developers are no longer able to bundle MySQL client support due to licensing problems with MySQL. The MySQL libraries are now GPL rather than LGPL, which makes them incompatible with the *Apache*-style licence of PHP. Work is under way by MySQL lawyers to make an exception for PHP to allow MySQL support to be bundled in future betas and the final release. In the meantime, you'll need to have your own client libraries installed if you want to try MySQL support in the beta.

Rethinking your algorithm

Write your own extension

If you've followed all the ideas given here (SQL-specific optimisations are coming next month!) and still not managed to get your code running at an acceptable pace, what do you do? Well, you're left with just two options: change your algorithm, or write your own extension. If you opt to write your own extension, you're in good company – just read through the PHP tutorials in the last three LXF issues to get all the help you need! When it comes to changing your algorithm, this is one of the most painful things a programmer has to do, but very often rethinking your plan is the only way to make real improvements.

```
ob_start();

echo "This is some text.<BR />";
echo "This is some text.<BR />";
echo "This is some text.<BR />";
echo "Last updated: " . date("H:i:s");

$content = ob_get_contents();
ob_end_clean();

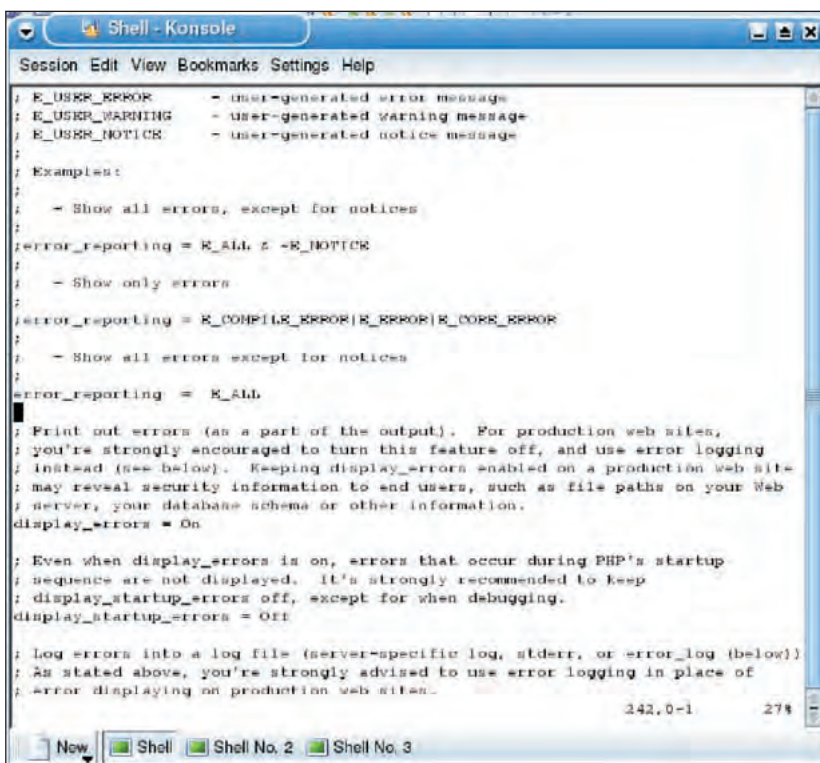
$handle = fopen("/var/www/public_html/$cachefile", "w");
fwrite($handle, $content);
fclose($handle);

include($cachefile);
?>
```

First, we use the **basename()** function, which strips out all path information about the **\$_SERVER['PHP_SELF']** variable, which holds the name of the current script. **basename()** takes a second parameter, a file extension, and if the filename passed as the first parameter has the file extension passed as the second parameter, **basename()** strips that off also. So, the first line will convert **/home/gallery/picture1.php** into just **picture1**, to which we append **.cache** to get our cache file. So, **\$cachefile** will be set to **picture1.cache**.

Then, **clearstatcache()** is called. PHP stores the result of all file checking functions in order to provide a speed boost, but as we're going to be checking file modification time regularly, it's best to clear this cache before we do any checks. The checks to perform are on the next line: first we check whether our cache file exists, and then we use **filemtime()** on **\$cachefile** to check whether the cache file is out of date. **filemtime()** returns its dates in a Unix timestamp, so we can compare that against the results of **time()** (the current time) minus 10 to get our 10-second cache life.

If the file exists and isn't out of date, the cache file is good to be re-used, so we just **include()** and exit the script as our job is done. If it isn't, it needs to be regenerated, and this is done using output buffering. Output buffering allows us to output content to a scratchpad, which we can choose to output to visitors or empty. **ob_start()** starts the buffering, and then we output four lines of text, including a time stamp so we can be sure caching is working. We then call **ob_get_contents()**, which returns all the content in our output buffer awaiting full output, and finally use **ob_end_clean()** to clear the buffer – the only copy of the four lines of text is now sitting in **\$content**.



To generate the cache file, we use **fopen()** to open the write file for writing, **fwrite()** out the contents of our output buffer, then close the handle. Finally, we use **include(\$cachefile)** again so that users always see the output as if we hadn't had to generate it from scratch.

Caching pages as simple as just four lines of text isn't going to help – in fact, it's likely to slow things down. But if you have a lot of work, such as reading the message list in a messageboard, it will help massively.

Use persistent connections

If you connect to a database with each script, consider using a persistent connection rather than a normal connection. For MySQL users, that's the difference between using **mysql_pconnect()** rather than **mysql_connect()**. Persistent connections remain connected even after your script has ended, which means that the next time a script asks for a connection, it uses the one that's already open – this saves a lot of time negotiating passwords and such that can otherwise be used to execute important code.

Switching to persistent connections doesn't require any other change than adding a **p** in the function name – the parameters are still the same. If your database server isn't on the same machine as your web server, consider using **CLIENT_COMPRESS** as the fifth parameter to your **mysql_connect()** or **mysql_pconnect()** call – it allows MySQL to compress data to save space, can drastically lower network bandwidth and transfer speed, particularly when reading in lots of data.

Experiment!

A large part of the optimisation process is very much based around a trial and error process – you see what works for your code in your situation, and try it out. Very often the best thing you can do is simply change the algorithm you're using – don't be afraid to try new things out! [LXF](#)

Set your **php.ini** file to **E_ALL**, so that it will show *all* errors in your scripts and never again be left in the dark.

NEXT MONTH

Next month we're going to be looking at turbo-charging your SQL queries by analysing queries, optimising table schemata, and using indexes.

K DESKTOP ENVIRONMENT

Beginners' Guide to Linux – KDE

PART 3 Offering a combination of familiarity and innovation, the K Desktop Environment is becoming the face of Linux for an increasing number of users. It could be all you need...

Linux is blessed with two complete, competing desktop environments which include everything from basic file management and information display to a growing range of integrated applications. While GNOME (which we'll be covering next issue) has been gaining favour with the likes of Red Hat and Ximian, developers such as SuSE, Lindows and Xandros have focused more on the K Desktop Environment (KDE), an ever expanding collection of applications, plugins and K-Parts developed to cover almost every area of personal computing.

There is a lot of detail to cover when it comes to KDE, but it all starts with the desktop and the Kicker/taskbar, the things you'll encounter when you boot up for the first time.

A. Desktop

It has been suggested that if you work with computers, you're likely to spend more time gazing at your monitor than at your loved ones, so it makes sense to at least make your PC's desktop look appealing! To quickly change the look of your desktop, right click on any clear space and select 'Configure Desktop...' from the menu. In the Desktop dialog, choose 'Background' and click on the 'Wallpaper' tab. There are quite a few options here, so go wild and experiment. Images can be stretched, tiled (for patterns) or centred and you can set up a common wallpaper, that is, it's the same on each of your virtual desktops, or choose a picture or pattern for individual desktops. KDE comes with a range of images for you to use, or you could hit the 'Browse' button to select your own pictures – perhaps of the loved ones you're missing while sat at your PC!

Fig1 The constituent parts that make up the KDE desktop. The letters refer to the explanation of each element in the text of this tutorial.


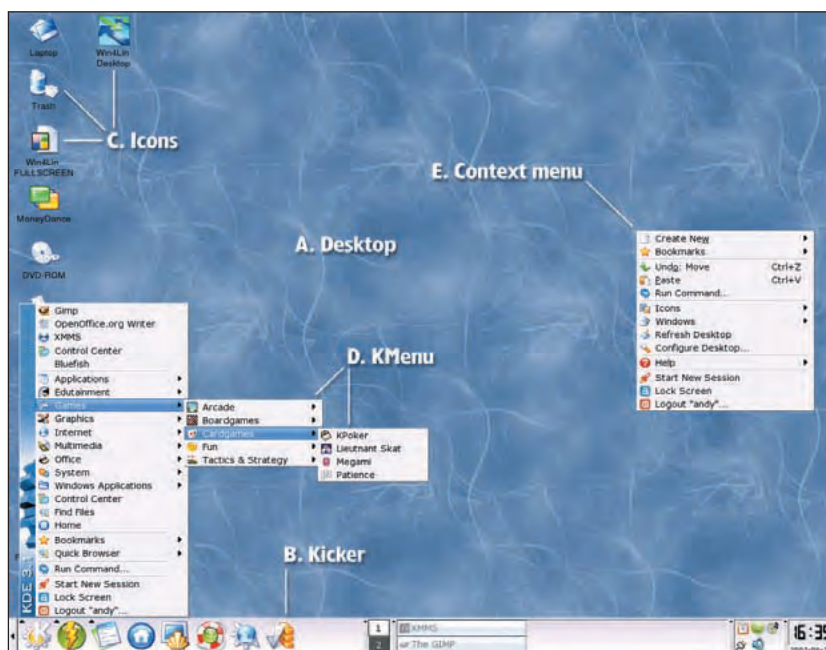


Fig2 The desktop background will load a picture, slideshow-style, from the selected directory every 45 minutes.

If you can't decide on a suitable picture, select Multiple wallpapers and choose 'Setup Multiple...'. Here you can choose to have a selection of images displayed, either in order or randomly, at intervals of your choosing. The Configure Desktop dialog box has many other options to play with, including screensavers, screen font selection, and virtual desktop setup as shown in Fig2.

B. Kicker

In KDE the task bar (usually running along the bottom of the screen, though this is not set in stone) is called *Kicker* and is divided in various elements.

The Kbutton on the far left of the bar is probably the first thing you'll click on. This is similar to the Start button in Windows and, depending on your distro choice, may have a 'K' icon, a lizard or a Red Hat on it. This is where you launch the *KMenu* (see point D on page 92) to open applications or access other parts of the computer such as the *Control Center*. Next to the start button, you'll often find some application launchers, a link to your /home directory and, occasionally, a document browser.

It's important to remember that, as Linux is all about choice, you aren't stuck with what you're given. For instance, you could add a new application to the bar by simply right clicking on a free space and selecting Add > Application Button > [selection]. The application, complete with icon, will be added to the bar. You can remove an icon by right clicking on it and selecting Remove [selection]. There is quite a range of buttons you can add here, including recent documents, *Konqueror* bookmarks, a quick file browser and *Desktop Access*, which handily minimises all the windows on a given desktop.

Next up is usually the virtual window manager. KDE can run a number of desktops which makes it easy to, for instance, have

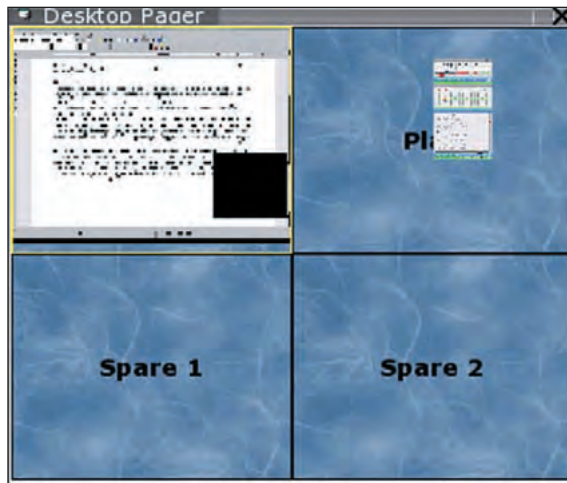


Fig3 The Pager gives you a visual idea of what's going on.

your background applications (email client, media player etc) running without their interfaces getting in the way of your main work apps. You can have anything up to 16 virtual desktops, but most people manage with two or four. To configure, right click on the 'applet' and select 'Configure Virtual desktops...'. This launches a dialog box which will let you choose how many desktops you require and also name them. If you want a more visual representation of what is going on in each individual desktop, right-click the applet and select 'Launch Pager' to see a display of all your desktops and their contents similar to that shown above in Fig3.

Beyond the Virtual Desktops is what KDE likes to call The Taskbar, a visual representation of running applications and open documents. Right-clicking on each individual bar will allow you to minimise, maximise, close, shade (just show the title bar) or send an entry to a different virtual desktop.

You can de-clutter your Kicker somewhat by moving the 'taskbar' section to the top of the screen where the buttons are still just a click away. Right-click on some free space and select Add > Extension > External Taskbar. By default, the new taskbar will appear below your existing Kicker, but you can change this by right-clicking the main Kicker again and selecting 'Configure Panel'. You should see something similar to Fig4. Now highlight 'External Taskbar' and select the Top Left icon in the position grid. Hit 'OK' and your task bar should be at the top of the screen. Finally right click the Kicker again and do 'Remove > Applet > Taskbar'.

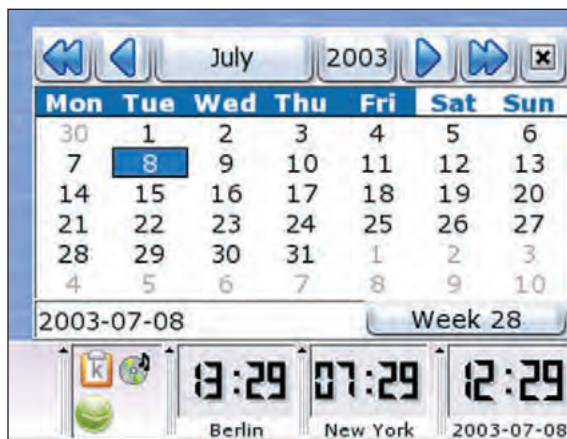


Fig5 Clocks and calendars galore.

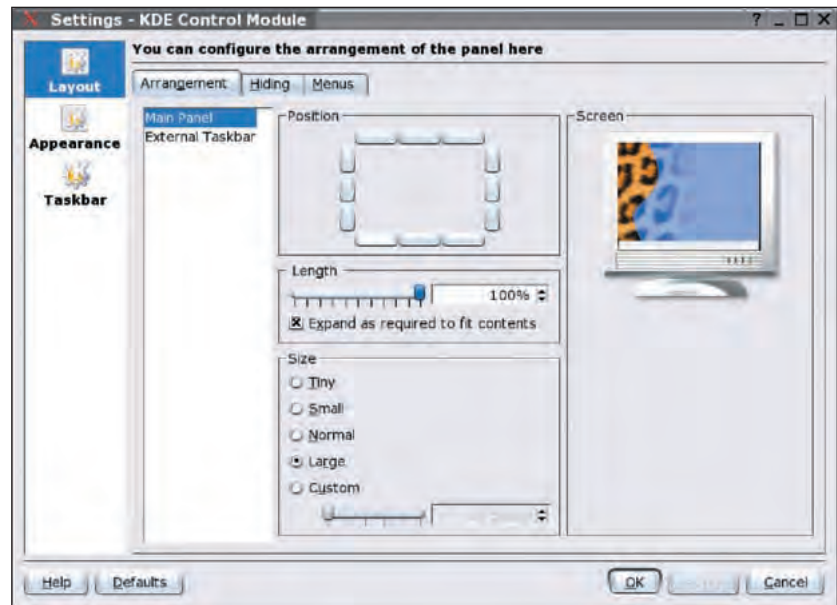


Fig4 The Kicker, and other taskbars, can go on any edge of the screen you choose.

Further to the right is the System Tray, which is a space used for quick access to applets such as the KDE Klipper, CD Controls and other useful gadgets. The Klipper stores a brief history of everything copied to the clipboard so you could, for instance, use it to store a series of cuts to be pasted into a document later on during a session. It can be configured to store as many entries as you like, though obviously if you're running a low resource machine, it does take up memory.

Finally we have the clock. But this isn't just any clock! Right clicking will give you a plethora of options including selecting different styles – analog, digital, plain or the always cool Fuzzy clock – a range of timezones, colours, time formats and other options (some would say too many). Left-clicking on the clock launches a small but useful calendar applet, shown on this page below left in Fig5.

The really great thing about the clock though is that you don't have to limit yourself to a single timepiece – especially handy if you're a commodities trader with an office in New York or a *Wolfenstein* player with an opponent in Berlin. To add a new clock to the Kicker, right-click on a blank space and select Add > Applet > Clock. The new clock will most likely be dropped onto the bar in the wrong place, so drag it to the correct location using the small grab handle to the left of the dial. You can then change the timezone by right clicking the clock face and selecting 'Show Timezone'. If your desired timezone isn't in the menu, select 'Configure Timezones' to add the one you want.

C. Icons

Icons on the desktop can be links to applications, files (images, MP3s etc.), websites, scripts or devices, and just like almost everything in Linux, there's plenty you can do to change them.

To see what we can do, right click on the CD/DVD icon and have a look through the menu. Some of the entries are obvious – 'Open' will mount the disk and launch Konqueror to display its contents, cut/paste works on the icon itself not the contents of the device – while others will be specific to the device.

One of the more important entries is 'Mount' which makes a volume (which can be a CD, DVD, Zip, USB camera, or any file system basically) available to the operating system. Most removable devices will be mounted as soon as you click an icon, but it's always good practice to unmount a device before you

Tip

If you've moved to Linux from Windows, you may notice that often icons are launched with a single click. While some people find this useful, it can be a chore if – being set in your ways – you keep accidentally starting two versions of everything!

Changing this to a double click is simple. Open the KMenu and select the 'Control Center'. Now click Peripherals > Mouse and then select the 'Double-click to open files and folders' radio button. Hit 'Apply'. Now single-clicking an icon or folder once will select it, which is good if you want to delete something, while double-clicking on an icon will launch it.

While you are in the Peripherals dialog, you could also select 'Keyboard' and change the Number Lock status at startup to 'Turn on'.

TutorialBeginners'Linux

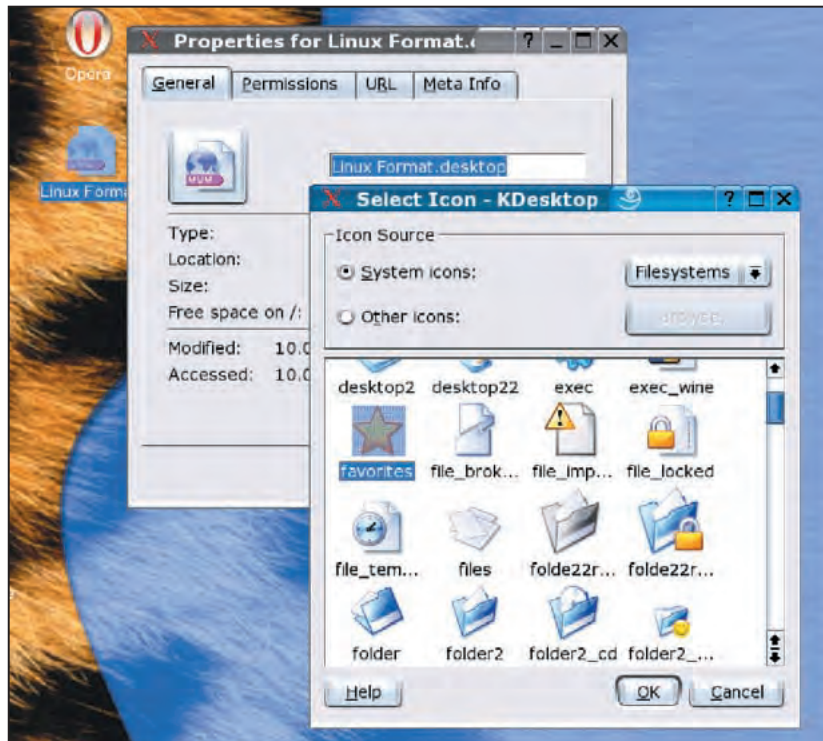


Fig6 KDE ships with plenty of icons for every conceivable occasion.

◀ remove the media by right clicking the icon and selecting 'Unmount'. You may find that, if you've just accessed a CD and then try to unmount it, you get an **unmount failed/device is busy** message. Don't panic! Give it a minute and unmount again.

You can attach lots of different things to icons. For instance, you could link to the *Linux Format* website directly from your desktop: right-click any free space on the desktop and select 'Create New > Link to Location (URL)'. Enter the name for the link that you are creating (*Linux Format*) and the web address (<http://www.linuxformat.co.uk>) and click 'OK'. A new icon with the default www image will appear on the desktop.

You can change this icon by right-clicking it and selecting 'Properties'. In the General tab, click on the icon image to launch a browser where you can select from a range of pictures split into groups such as Applications, File Systems and Mimetypes. If changing one icon at a time doesn't satisfy, you can change them

all by opening the Control Panel and selecting Appearance and Themes > Icons and choosing from the installed sets. More can be downloaded from www.kde-look.org and can be installed from the original zip file. Just select Install New Theme and browse to the download location and hit OK. The icon set will then be available in the main dialog box.

D. KMenu

You may have noticed by now that the KDE project is very fond of the letter K – it's the one before 'L' for Linux. The *KMenu* is where you access applications, help, documents and system settings. It looks quite bewildering, especially if you've opted to install millions of applications, but on most distros the structure is quite logical; have a browse.

Don't like where things are? Well, right-click on the *Kicker*, select 'Configure Panel...' and the 'Menus' tab. Here you can do big changes such as adding a Recent Documents entry to the main menu, or, by hitting 'Edit K Menu', you can burrow into the actual menu and add, remove or alter individual entries as in Fig7.

E. Context menu

This is what appears if you right-click on a bare space on the

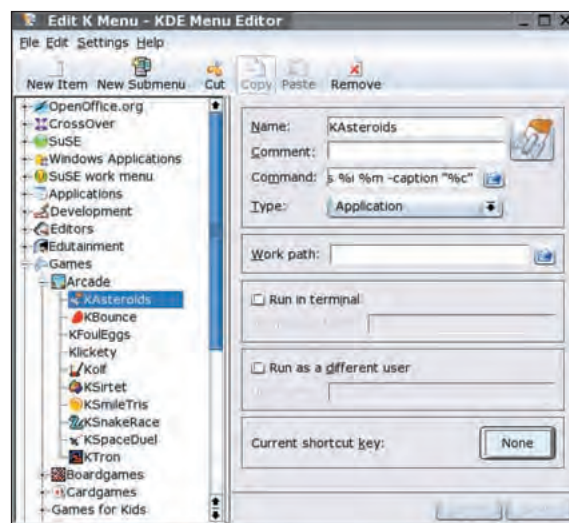


Fig7 Don't like the menu setup? Then change it.

What on Earth is Karamba?

Dynamic system information


AY KARAMBA! ANYTHING TO DO WITH THE SIMPSONS?

Nope. *Karamba* is an extension to KDE which, among other things, allows KDE to have its own MacOS-X style dock with zooming icons and all manner of eye candy. It was inspired by a Windows project called *Samurize*

(www.samurize.com) which started life as a tool for adding dynamic system information to the Windows desktop.

AND WHAT IS THE POINT OF IT?

As well as beautifying your desktop – if you like that sort of thing – it can also be used to monitor various processes including network access, memory/hard disk usage, email and the weather.



However, the most popular use seems to be as a replacement for KDE's *Kicker*. **CAN'T YOU DO ALL THIS STUFF WITH KDE ANYWAY?**

Yes you can, but *Karamba* offers more flexible themeing using Python scripts. The disadvantage though is that as *Karamba* is a very young project (just a

couple of months old in fact, at time of writing) it's aimed at the user comfortable with the concept of editing the coordinates in a text file to make the Mac OS X style bar (or whatever) sit on the bottom of the screen properly. **CAN I SEE WHAT IT'S ALL ABOUT?**

There are a couple of useful places to

visit on the Internet. First up is www.kde-look.org where *Karamba* has its own subsection housing projects, forums and screenshots.

The other worthwhile link is the home page of Hans Karlsson, who wrote the original script:

<http://www.efd.lth.se/~d98hk/karamba/>

desktop. The same mouse action on an icon, folder or the Kicker will bring up a different menu.

Eye candy

Making your desktop look nice is not just good for your eyes. If you – or your users – are used to Windows or Mac OS, for instance, you can mould the KDE desktop to mimic that OS, with similar window decoration, style and colours. You control the visual aspects of KDE by launching the Control Center and navigating to the 'Appearance and Themes' section as in **Fig8**.

We've already covered the icons and backgrounds earlier in this tutorial, but there are many other things you can change in here. The two most significant elements are Style and Window Decorations. The former determines how UI widgets, that is buttons, tabs, progress bars and the like, look. If the chosen style supports it, this is also where you set menu effects such as transparency, animation and fades, and also where you can alter a selection of toolbar settings, add text to buttons or stick the application menu to the top of the screen Apple style.

KDE's default style is Keramik, which is blessed with a number of good-looking effects which should only be attempted on a processor faster than 500MHz. If you have processing power to spare, select the 'Effects' tab and click on 'Enable GUI effects'. Click on the 'Menu effect' drop down and select 'Make Translucent'. You can also add a drop shadow with the radio button below and adjust the opacity of the menus. Hit 'Apply' and open the *KMenu* to see the effect. Don't go overboard though, the idea behind a GUI is to make your computer easier to use, not beautify it to death, as it has been in **Fig9**!

The Keramik Style is complimented by a similarly named window decoration, though the options available here aren't quite as extensive. Once you've selected the window decoration, you can go into the 'Colors' section and either choose one of preinstalled colour schemes or create your own. Keramik comes with a trio of schemes and there are many more installed with KDE. If one of the schemes is almost what you want, you can alter specific parts of it by selecting what you want to change – desktop background, active title bar, etc – and then selecting the new colour. Remember to hit the 'Save Scheme' button and give your work a name!

The great thing here is that you can't mess up your system (unless you make everything black) and if you end up with something that would make even Lawrence Llewelyn Bowen wince, you can just move through the menus hitting 'Default' to get back to the appearance that you started with.

Konqueror

At the core of KDE is *Konqueror*, an application for which the epithet file manager is a definite understatement. It is, thanks to the modular design of KDE, equally at home dealing with local file systems, web pages, photo albums, sound files, ftp locations and a variety of office formats. We'll look at using *Konqueror* as a web/network browser, and as a multimedia organiser later in this series, so we'll just mention basic file management and previewing here.

The quickest way to get to *Konqueror* is by clicking the Home icon on the *Kicker* or desktop. This will launch the application within your /home directory, which is the best place to save your documents, photos and music files.

Just below the top menu bar is the toolbar which is used for basic navigation. The up arrow will move you up the structure of your file system until you get to /root, while the back and forward

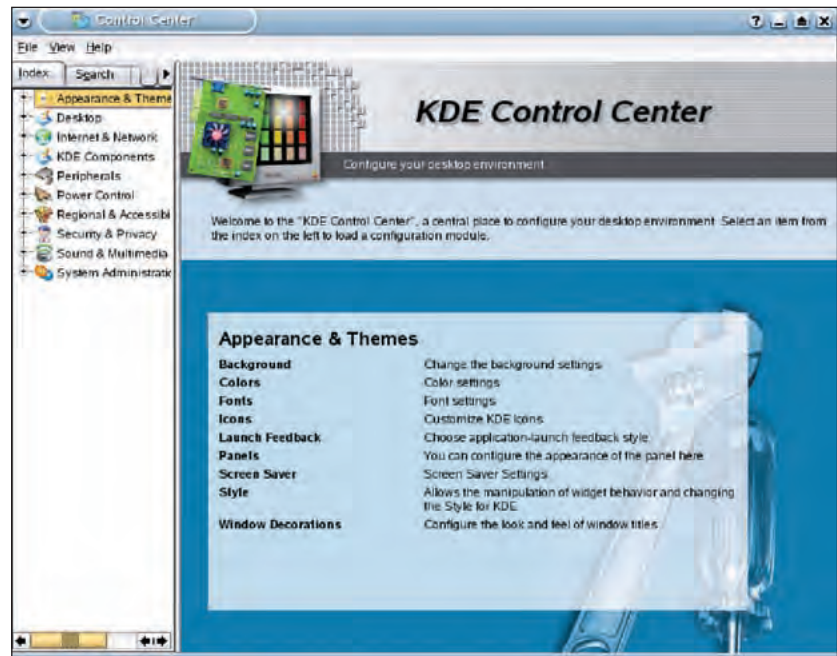


Fig8 Change your themes in *Control Center* – the computer equivalent of a make-up bag.

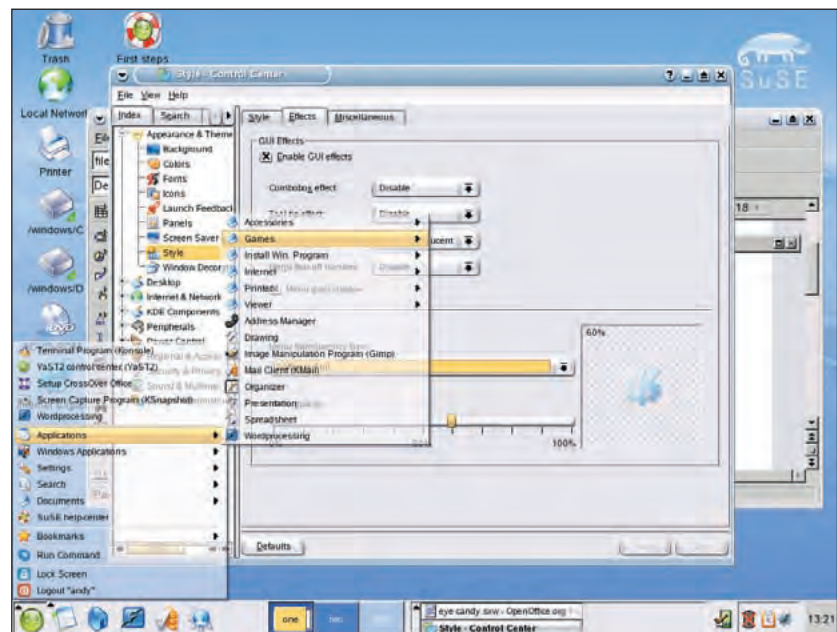


Fig9 Take it easy with options like translucent menus – keeping your system useable is very important!

buttons allow you to navigate to previously viewed folders. You'll also see icons for home, refresh, cut, copy, paste and print – which all do as you'd expect – and two magnifying glasses to alter the size of icons or preview images.

The simplest way to copy or move a file is to open two *Konqueror* windows – one at the source, the other at the destination – and then drag the file across. When you drop it in the window, a dialog will ask whether you want to move or copy the file. You could also use the cut/copy and paste method, or right-click on the file and select 'copy to/move to' and use the quick browser to select the destination.

The right-click, context sensitive menu can also be used to create new folders or documents, open documents in specific applications and preview files. [LXF](#)

NEXT MONTH

We've not finished with KDE or *Konqueror* (especially its multimedia features) yet, but next time we'll be looking GNOME, the alternative desktop environment. Before then, see if you can get *Konqueror* to show, play and rip (to MP3 or Ogg) the contents of an audio CD without using any other application.

Answers

If you are really stuck and the HOWTOs yield no good result, why not write in? Our resident experts will answer even your most complicated problems!

Our experts

Whatever your question is, we can find an expert to answer it – from installation and modem woes to network administrations, we can find the answer for you – just fire off a letter or email and it'll all be taken care of.


LXF answers guy **David Coulson** is a networking and security guru with plenty of sysadmin experience to boot.



Nick Veitch is the editor of the magazine, and answers your easy questions! Or indeed anything to do with *Grub*, *LILO*, *netatalk*, vi...



Hans Huberland is Rackspace Managed Hosting's Linux expert. Send any Linux system admin questions to sysadminqa@rackspace.co.uk



Bye, Windows!

Q I am about to install Mandrake 9.1 as given away in your June edition. I just need a little bit of advice.

I run Win2000 and WinME on a 20GB Hard drive. One OS must go, but which one? It does not really matter which – I am open to suggestions. Finally, should Mandrake be installed on my computer's 'C' or 'D' drive?

Ronald Wood, via email


A As to which Windows operating systems should bite the dust – we couldn't decide, so just get rid of them both! Actually, since Win2k is based on NT and WinME is based on Win95, you might as well just drop WinME from the system, as Win2k is vastly more stable, although that isn't saying a whole lot.

Mandrake has no concept of 'C' or 'D' drives, as this is the Windows way of giving useful names to media devices. You can install Mandrake on either device, although Windows does not like being moved from one device to another, or having the D: drive becoming C: if Windows is already installed on D:; due to items in the registry being incorrect.

Linux modems

Q Please please can you tell me a modem that works on Linux. Every dealer I've asked looks at their modems and as it's not stated on the box says it *might* work? All the articles I've read, including your *Complete Guide To Linux* says it's hard to get a modem to work, but no one it seems can give a make and model that when I buy it, it will work.

So please could you send me the details of a PCI Modem that works? An article in your great mag would help a lot of newbies over this barrier I'm sure.



Linux does not use 'C' or 'D' drives, instead defining filesystems by the device and partition number.

Thanks for the DVDs; Mandrake 9.1 is great: picking every bit of hardware, Plustek Scanner and Olympus Digital camera, but the modem – no! I don't understand it all but keep up the great work in your mag I'll get there in the end.

Gerry Green, via email

A The simple way to get a modem that will work 100% with Linux is to get an external serial modem. Any internal PCI modem which you wish to use with Linux will need to be a hardware modem, in that it actually has a serial device, otherwise known as a UART, on board. If you can look at the modem in a store to check, you'll probably be able to find the UART on it. We don't have any specific brands or models to recommend, although with the recent increase in broadband usage across the country, eBay may be a good location to find a cheap modem that will work with Linux.

Powercut problems

Q I am running Mandrake 9.1, installed from the June LXF DVD. All OK until I had a power cut while the system was running. The next boot seemed to check/fix the filesystems OK and the system came up and

subsequently shutdown, with no obvious problems. However on subsequent attempts to boot into Mandrake I get the following messages, and the system then hangs:

Uncompressing Linux... OK Booting the Kernel

Kernel panic: No init found. Try passing init= option to Kernel

Any ideas on what has happened, and what *init* option I need to specify? I have 3 systems on the disk (Windows, Red Hat 7.2 and Mandrake 9.1) and boot from a GRUB floppy. Windows and Red Hat both boot OK.

Les, via email

A When the kernel gets started up, it will attempt to start the *init* process by executing `/sbin/init` from the root filesystem. It would appear that your system can't find `/sbin/init` on your root filesystem. You could try to pass `init=/bin/sh` on the command line to get a shell prompt then look in `/lost+found` to see if `/sbin/init` ended up in there following the fsck.

Mandrake 9.1 also has a 'rescue' mode where you can boot off the CD and mount your root filesystem to fix any problems. If `/sbin/init` has disappeared, then you will have to

reinstall the RPM package from the CD which provides this file to replace it


Drives me crazy

Q I have a question regarding attachment and configuration of SCSI tape drives.

I am trying to connect a HP Surestore 6000, 4mm tape drive to a system running Red Hat 8. The SCSI card is being seen by the OS at boot time, however, KUDZU is not seeing the tape drive (well, I figure that is what should be used to configure it successfully). The BIOS on the SCSI card is seeing the tape drive as well. I have tried using the tape drive on a Windows 2000 system and it configures OK, so it's not a tape drive issue (unless the tape drive is not compatible with Linux of course).

Could you possibly point me in the right direction for fixing this one? I have been searching for documentation on adding devices to Linux, but I am either not looking in the right places or there just isn't that much out there (and I can't really see that being true!) Your help would be greatly appreciated.

Nigel Henderson, via email



/proc/scsi/scsi shows information on all the attached SCSI devices, although for hot-pluggable devices it needs a kick with **rescan-scsi-bus.sh**

The first step is to check `/proc/scsi/scsi` to see what devices the kernel has detected. **dmesg** may also come in useful, as this will log when a device has been located by the kernel. If the kernel has support for the SCSI controller, you may need to manually **modprobe** for the kernel module for it to initialise the controller.

Unfortunately you did not provide us with any details as to what your

SCSI controller is, nor if the Linux kernel does anything with it when you boot up. You also need support in the kernel for SCSI tape devices, which may be compiled as a module called 'st.o'. Once this is loaded and the kernel has detected your tape drive, you will be able to use /dev/st0 and /dev/nst0, which is the non-rewinding version, with a variety of tape backup tools. If you still have no success, if you are able to send *LXF* the output

from **dmesg** as well as more hardware details for your SCSI system, we may be better prepared to help you further.

MDK modprobe II

Q In response to “MDK modprobe”, *LXF* July 2003; I, like most Red Hat users I

I, like most Red Hat users I would imagine, share Mr Ewing's grief when it comes to unsatisfied RPM dependencies. The amount of free software available for Linux is amazing, and precompiled binaries can make software installation really quick and easy, but the number of times I have given up half-way through an ever-growing maze of dependencies is countless!

Whenever I moan about this, I get one of two responses; either a shrugging of shoulders and an understanding frown, or just “use Debian, it’s 133t!”. Whilst I will definitely give Debian a bash as soon as I have my own broadband connection to download the “testing” version, nothing can beat the simplicity of bunging in a couple of Mandrake CDs and getting all the shiny new versions of everything. My webservers also run a Red Hat derivative, and I don’t feel

A QUICK REFERENCE TO: **Kernel compilation**

Generally once Linux has been installed, many people like to head over to ftp.kernel.org, and download the latest kernel release. This ensures that they are running a system with the latest bug fixed, which may affect stability and security, as well as providing any additional device support for specific hardware which their system has. Many distributions come with kernel sources, which make it easy to add or remove support for certain things, but if you are wanting to apply third-party patches to the kernel, or otherwise change it, you will need to ensure that you have a version of the kernel which is clean, otherwise the patches will fail.

Rebuilding the kernel has a number of stages, which involve setting up the kernel with the appropriate capabilities, compiling the code for the kernel, and then installing both the kernel, and any modules, on the system. The

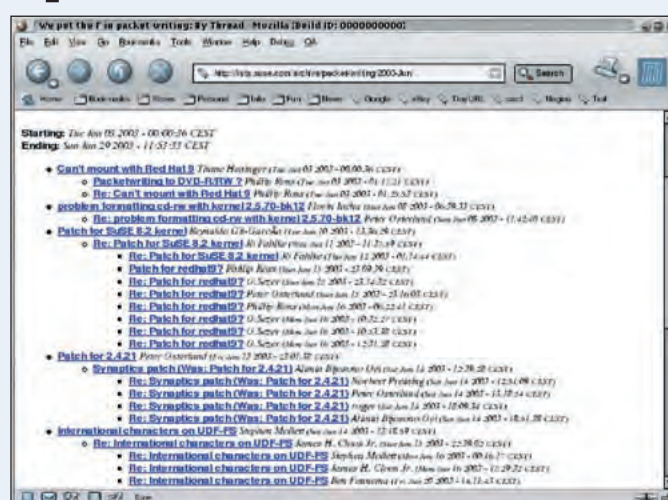
following stages will work equally well with either a kernel downloaded from ftp.kernel.org, or the source code for a distribution's kernel (generally found on the CD, or on the ftp site).

The kernel configuration is stored in `linux/.config`, which can be copied out and saved for use in a later kernel. If we store our kernel's configuration elsewhere, we can copy it back into our kernel tree, as `linux/.config`, then run:

```
# make oldconfig
```

which will step through the existing `.config`, and will prompt us for any new configuration options that will have been included since we last built our kernel. This equally works if we patch our existing kernel, and want to know what new things it can do.

If we want to reconfigure our kernel, we can use **make config**, **make menuconfig**, or **make xconfig**, which use different user interfaces for configuration. **menuconfig** should be used for a console system, or **xconfig** for a box with an X server running.



Packet writing support for the Linux kernel.

Once our kernel is configured, we can build and install it in one go;

```
# make dep && make clean &&  
make bzilzo && make modules &&  
make modules install
```

This will clean up the kernel tree, removing old binaries, build the new

one, then install it with `lilo` and install the modules. **make bzlilo** installs the kernel in `/vmlinuz`, and runs `/sbin/lilo`, so if you have your kernel stored in `/boot`, you will need to copy it over manually and rerun `/sbin/lilo` manually.

FREQUENTLY ASKED QUESTIONS **KERNEL MODULES**

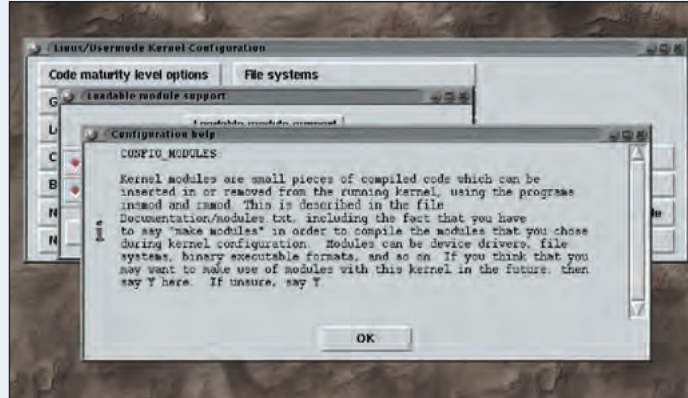
FAQ WHY WOULD I WANT TO USE KERNEL MODULES?

Kernel modules are great if we want to save some memory by not having support for every piece of hardware we own loaded in the kernel whenever the system is up and running. If we only use certain filesystems occasionally, or have a number of USB devices that may or may not be connected, we can save a few hundred kilobytes of RAM by having them as loadable modules.

Modules are also particularly useful if a device needs some options passed to the kernel module with an IRQ, IO, or other specific information. Rather than rebooting the system to try different options as you might usually expect, using a loadable module allows various values to be tried out without having to restart the kernel.

FAQ IN WHAT SITUATIONS WOULD I WANT TO AVOID MODULES?

Generally one would not want to have loadable modules on a server, or a system where there is really no benefit from having kernel modules. If the hardware rarely changes, or it is important for all of the devices to



Linux Kernel Module support allows support for specific devices to be loaded and unloaded on the fly without having to reboot.

function with the bare minimum of configuration, then compiling it all into the kernel is a must. If the root filesystem becomes unreadable, then you can't load the required modules.

It's also often important to remove loadable module support where security is a concern, as a number of rootkits include a kernel module which is used to rewrite some of the system calls so you can't find it's processes or files on the local system. However, even without Linux Kernel Module support, it is still possible to load an exploit via /dev/kmem, although patches exist to block writing to this device.

FAQ HOW DO I ACTUALLY COMPILE A LINUX KERNEL MODULE?

Kernel options are selected to be compiled as modules within the **make menuconfig** or **make xconfig** kernel configuration tools. To select a device as a module 'Loadable module support' must be selected and optionally the kernel module load (*kmod*) which will automatically load filesystem and device modules when specific devices are accessed.

The kernel can be compiled in the regular way with **make bzImage** or **make bzImage**, but the modules

must then be compiled with **make modules**. Once the modules have finished compiling, they must be installed with **make modules_install** which will install them into /lib/modules and use **depmod** to build a module dependency tree.

One must ensure that the latest version of *modutils* is installed on the system, as there were a number of changes to the organisation of /lib/modules in the 2.4 kernels.

FAQ WILL THE KERNEL LOAD THE MODULES AUTOMATICALLY, OR DO I HAVE TO DO IT MANUALLY?

With *kmod* support, the kernel will load modules for filesystems and some devices on the fly without any interaction required. However, in many cases, the modules will not be loaded automatically. Depending upon the device they are used with, one can setup automatic module loading via *devfsd*, *modules.conf* or in the case of USB and Firewire devices, *hotplug*.

devfsd is a management daemon for users of the *devfs* filesystem, which calls a module load each time you open a device. For example, if you try to open /dev/net/tun and it

« confident enough to do a remote Debian install!

So what are the options?

Jon Melhuish, via email

A Red Hat seems to have not helped matters any with its latest effort, as many people have given up with Red Hat following 9.0 due to dependency issues. There are many more useful distros when it comes to maintaining dependencies, including Debian and Gentoo. The latter requires packages to be compiled locally, which becomes rather impractical on a server where CPU cycles are better used handling DB queries or mail delivery than recompiling *glibc* and *tex*.

There are frequent Debian releases, the most recent of which is 3.0r0, which can be obtained on CD from the usual locations. This is Debian stable, rather than testing, although it's very straightforward to dist-upgrade a Debian box from stable

to testing. As Debian is somewhat more difficult to install and maintain than Mandrake and Red Hat due to the lack of GUI tools, it may be a wise idea to have a friend with broadband download the ISOs for you and burn them to some CDs, or buy them from Cheapbytes or Linux Exporium for a few pounds. This way, you can become used to the way Debian works, and forgo downloading packages over a 56K modem.

Another alternative is to use 'User-mode Linux' to run two distributions at the same time. The existing distribution would continue to be installed and maintained, but one can run other distributions as user-space processes. This is great for playing with other distributions without having to reinstall. User-mode Linux can be found at <http://user-mode-linux.sf.net/>. For more on Red Hat package management, see page 62.

Packet Writing

Q Fortunately with *Linux Format* magazine, I learn more and more every month. No other magazine that I have found provides as much information, product reviews and tips (and don't forget the software on CDs) than what is found within your pages.

There is one main thing that I still can do in Windows which I haven't figured out how to do in Linux. Within Windows I can quickly and easily read and write to my CD-RW disks as if they are a floppy drive. This is very handy to move files back and from between my office and home PCs. How can I do this in Linux? This one problem, (besides some killer games, of course!) is keeping me from becoming totally Linux.

Bob Ackerman, via email

A What you're talking about is packet writing, which allows packets of data to be written to a CD-R/W device as if it was a regular block device. To do this with Linux you need a fairly untested kernel patch, which can be downloaded from <http://w1.894.telia.com/~u89404340/patches/packet/2.4/packet-2.4.21.patch.bz2>. This will require the latest release of the Linux kernel and probably won't patch properly with vendor kernels, such as those from Red Hat and Mandrake, so if you're unfamiliar with compiling your own kernel you may wish to give this a miss and wait for it to be included by vendors if it has not already.

SuSE actually has a number of mailing lists for those interested in packet writing support for Linux, although they are somewhat technical. You can join at this address:

does not exist, `devfsd` will do `modprobe /dev/net/tun`. By creating an alias in `modules.conf` for `/dev/net/tun`, we can have the `tun.o` module loaded when we access `/dev/net/tun`.

```
alias /dev/net/tun tun
```

If `devfs` is not being used, we can load modules based on the 'major' number of the device we access in `/dev`. For example, if we want it to load the `tun` module automatically, we need to find out the major of `/dev/net/tun` which is done with `ls`:

```
macha:~# ls -l /dev/net/tun
crw-r----- 1 root root 10,
200 Dec 31 1969 /dev/net/tun
```

We can see that this is a **char** device by the first **c** in the output from `ls`, with a major of 10. When we access this device the kernel will try to load **char-major-10**, so we can simply alias this to `tun` and it will load `tun.o` for us.

```
alias char-major-10 tun
```

FAQ CAN I LOAD SPECIFIC MODULES WHEN I BOOT UP THE SYSTEM?

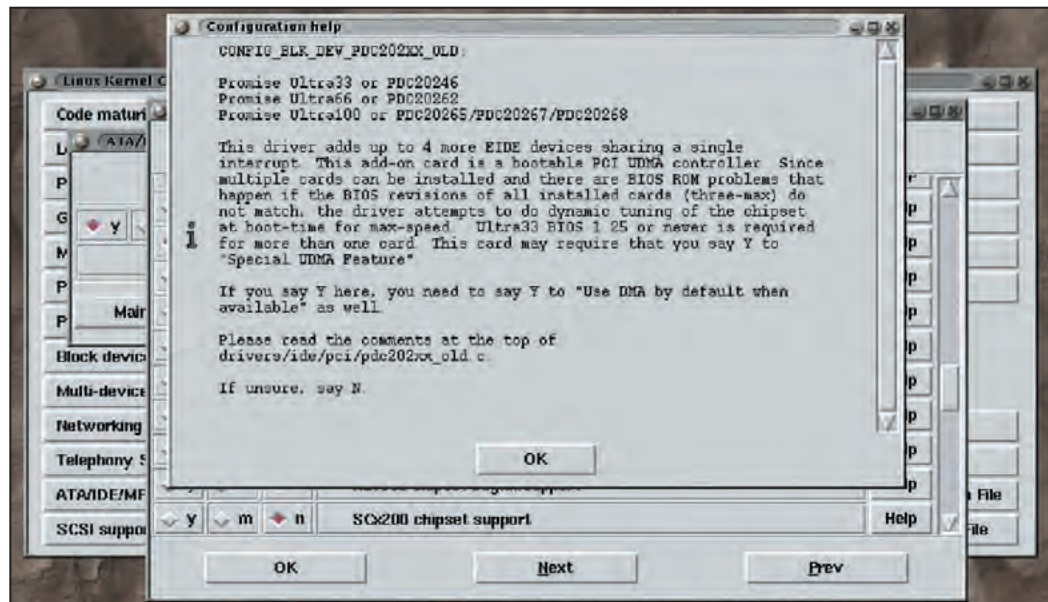
Depending upon the distribution, one can either put `modprobe` commands in `/etc/rc.d/rc.local`, or into `/etc/modules`. On Debian systems, `/etc/modules` is used and a simple list of all of the modules to be loaded at boot time is included in this file.

www.suse.com/us/private/support/maillinglists/

Retrograde MDK

Q I am running Mandrake 9.0 on a Pentium III, 768MB RAM and 100GB hard disk – not all of this space is available to Mandrake, as I also run Win2K. The disk is mounted on Promise FastTrak 100 RAID controller. Mandrake does not claim to support this, but it is working. At the weekend I tried upgrading to Mandrake 9.1 using your cover DVD. This failed as it could not locate a hard disk partition for the installation – my guess is that it no longer recognises the controller. I know that I should convert to Red Hat or SUSE (both of which support the controller) but I am curious as to why I can run MDK 9.0 but not 9.1.

John Huckle, via email



The Promise FastTrak controller in Mandrake 9.1 doesn't play nicely, although a kernel upgrade will fix it.

A Strange that Mandrake 9.1 does not include support for the Promise RAID controller. However, this does appear to be an acknowledged issue in its 9.1 release, as it there is a bug report for it at http://qa.mandrakesoft.com/show_bug.cgi?id=1268

Your best option at this point is to go back to 9.0 and wait for them to release a replacement ISO for 9.1 with the appropriate fixes, or for 9.2 and hope that they managed to fix this. The underlying reason for it failing is probably due to kernel changes between the releases, as the controller is recognised by no devices on it are detected by the kernel. There may have been someone on the Mandrake mailing lists who built a new install image for the distribution using a kernel without this problem, although one may have to resort to Google or requesting information on Mandrake mailing lists to locate such an image.

MDK9.1 & monitors

Q I have been unable to install Mandrake 9.1 from the 2 coverdiscs which came with the June edition of LXF. I am running Mandrake 9.0 which I bought from the Linux Emporium about 6 months ago. I have tried both upgrade and clean install in separate 6GB partitions. (I have Partition Magic 8) The installation goes fine up to the point where I configure the graphical interface. The installer chooses exactly the

same parameters for 9.1 as are used by 9.0. That is: SiS630, generic monitor 1024 x 768 @ 70Hz with 800 x 600 at 24bits, then the program hangs at the test screen. I have had about six attempts. My monitor is a Maxdata Belinea 10 70 20 of size 39.6cm diagonal which works fine with 9.0. My questions:

1 Is there something I can do to fix the problem when I use your coverdiscs?

2 If I get 9.1 from somewhere else, will I still have the same problem? I am not writing so as to blame anyone, just to fix the problem – if a 3 CD set of install discs will do the trick then I will get hold of some. This, after all, is the one of the beauties of Linux, you can afford to upgrade regularly.

Richard, Wolverhampton

A This error is down to an update in the XFree86 version in the Mandrake 9.1 release. As the Mandrake 9.1 files which are distributed on the LXF cover discs are exactly the same as that you can download from their FTP server, you'll not have much success obtaining it from elsewhere. As it locks up, it's likely to be a specific problem with the utility Mandrake uses to probe for a AGP or PCI video card and unfortunately there isn't really very much you can do about this. If there is a way to explicitly specify your video configuration and avoid having it probe, that may solve the problem.

You may want to backup your `/etc/X11/XF86Config` file from your

existing Mandrake 9.0 installation and then upgrade it to Mandrake 9.1 using a simple VGA video card configuration. Assuming this solves the problem, you can then copy the previous `XF86Config` file over the one which Mandrake 9.1 created and restart X so that it uses the alternative configuration.

Submission advice

We are happy to answer all sorts of Linux related questions. If we don't know the answer, we'll find out for you! But in order to give you the best service, it helps a lot if you read the following submission advice.

- Please be sure to include any relevant details of your system. 'I can't get X to work' doesn't really mean anything to us if we don't know things like what version of X you are trying to run, what hardware you are running on.
- Be specific about your problem. Things like 'it doesn't work' or 'I get an error' aren't all that helpful. In what way does something not work? What were you expecting to happen? What does the error message actually say?
- Please remember that the people who write this magazine are NOT the authors or developers of Linux, any particular package or distro. Sometimes the people responsible for software have more information available on websites etc. Try reading the documentation!

We will try and answer all questions. If we don't answer yours specifically, you'll probably find we've answered one just like it. We can't really give personal replies to all your questions.

WRITE TO US AT:
Linux Format, Future Publishing, 30 Monmouth Street, Bath BA1 2BW or email: lxf.answers@futurenet.co.uk

Answers



Weakly message!

Q I have a server which runs the *Plesk* control panel interface which in turn uses *Qmail*. I've recently had a complaint from one of my users that it takes a week for the server to notify him of a message that is not delivered. Below is the content of the mail as returned by *Qmail*.

Hi. This is the qmail-send program at domain.com. I'm afraid I wasn't able to deliver your message to the following addresses. This is a permanent error; I've given up. Sorry it didn't work out.

<user@place.com>:

> Sorry, I couldn't find any host by that name. (#4.1.2) I'm not going to try again; this message has been in the queue too long.

Is there any way to tweak this setting to something lower? Thanks!

Roberto, via email

A *Plesk* does not allow much direct control of any of the software that it runs on. The reason for this is that it keeps almost all of its settings in a MySQL database which it uses to regenerate the config files for all the daemons it depends on. However, I think you may be in luck with the particular setting you want to change – the file does not exist by default so I'm fairly certain that *Plesk* will not interfere if you create the file.

By default *Qmail* tries to deliver mail for 1 week (40 times). Its retry schedule gradually increases from a few minutes between retries to about a day between retries at the end of the week. There is a single flat file by the name of *queuelifetime* under your *qmail/control* directory. By default this file does not exist, so *Qmail* uses its default of 604800 seconds (1 week). If you drop the number down to something lower like 176400 (22 attempts) mail will only live in the queue for about 2 days.

Plan partitions

Q Please help me with a problem I've been having. I have a system running Red

Hat 8, which has been reloaded before – because I underestimated the importance of planning my partitions. This time round I left plenty of space on one of the two hard disks to make new mount points, as I need them. That time has now come and I'm trying to add a mount point for */var/log* to keep my ever growing logs files together in their own location. After a reboot I find that */home* is empty – as you can see below this directory has its own partition and I have not touched it at all. Here is my */etc/fstab* in case that helps you:

```

LABEL=/ / ext3
defaults 1 1
LABEL=/boot /boot ext3
defaults 1 2
LABEL=/home /home ext3
defaults 1 2
LABEL=/var /var ext3
defaults 1 2
LABEL=/data /data ext3
defaults 1 2
none /dev/pts devpts
gid=5,mode=620 0 0
none /proc proc

```

```

defaults 0 0
none /dev/shm tmpfs
defaults 0 0
/dev/hda2 swap swap
defaults 0 0
/dev/hdb2/var/log ext3
defaults 1 2

```

I can't see anything obviously incorrect myself. Is there anything I may have missed?

Ant, via email

A At first glance I cannot see anything unusual with your configuration. I do notice though that you've using labels to identify your original mountpoints and a device name for the new disk. What is the label on */dev/hdb2*? If you created the partition on *hdb* in the exact same start point where there was previously a labelled partition, that new partition will adopt the label of the old partition. Partitioning the drive will not erase any data and formatting wont get rid of the label, as it is not a part of the filesystem's data area but written to the very front of the actual partition with other information. Check the output of:

```
e2label /dev/hdb2
```

and I think you may find the source of your problem.

In fact, I would personally recommend not using labels to reference mount points unless you find yourself frequently changing disk positions around on your IDE channels. A single messed up label

can prevent you from booting the entire system.

Dependency hell

Q I really love the *LXF* Answers section because I get to see real-life problems with interesting solutions. I hope you can provide an interesting answer to my question. Being fairly new to Linux I've fallen into what I believe is referred to as RPM Hell. When I try to install a package it sometimes complains that it needs another package, this is pretty straightforward if you have the package and are ready to install it. Other times it complains about a missing file or library and there is no RPM by that name available. I know I can search on www.rpmfind.net but it always brings up a lot files for other distributions, versions or architectures. Is there no way to make rpm report which rpm package the files it's looking for are meant to be in? Some hardcore friends have recommended I use Debian but I'm a little frightened.

Susan, via email

A Hello Susan – fear not! If you're using a fairly new version of Red Hat you will find an rpm for a package called *rpmdb-redhat-<version>-<date>* on the CDROM. If you install this you'll be able to use two useful new

About Rackspace Managed Hosting

Rackspace Managed Hosting™ provides complete managed hosting solutions with Fanatical Support to small to medium-sized customers as well as large enterprises. All customer platforms include state-of-the-art data centres,

customised servers, burstable connectivity, 99.999% uptime SLA, a dedicated account manager, instant emergency response and access to live expert technicians 24x7 for support of all hardware and core software. Founded

in 1998 with locations in San Antonio, Texas and London, the company manages over 7,000 servers for customers in more than 80 countries. For more information, visit www.rackspace.co.uk and see this month's *Linux Pro*.

The screenshot shows the Rackspace Managed Hosting website. The header includes the Rackspace logo and navigation links. The main content area features a large banner with the text "It's not about SERVERS, it's about SUPPORT." and a photo of a smiling woman. Below the banner are several sections: "Fanatical Support™", "Did You Say 100%?", "Happy Birthday Rackspace", and "Move to Rackspace and Save". The "Did You Say 100%?" section highlights a 100% Network Uptime guarantee. The "Happy Birthday Rackspace" section celebrates the company's 4th anniversary. The "Move to Rackspace and Save" section mentions potential savings of up to \$50,000 annually. The footer contains contact information and a search bar.

WIN A SHARP ZAURUS PDA FROM RACKSPACE MANAGED HOSTING

www.rackspace.co.uk

Every month, the best question related to Systems Administration that a LXF reader sends in wins a Sharp Zaurus SL-5500 running Linux.

With industry support from the likes of IBM, SAP, Oracle, and Sybase, the Zaurus features a QWERTY keyboard behind a sliding cover. The 65,536 colour screen is 240x320 resolution. An Intel® StrongARM® 206MHz processor and 64MB memory makes the Zaurus ideal for downloading software or playing MP3 and MPEG1. The docking station that fits into its USB connector, allowing users to synch with their desktop, schedules, address

book and ToDo lists. The email app supports POP3 or IMAP email servers. Write and send emails and browse the Internet with a compatible modem card, or using the infra-red port and a mobile phone; or any replied, forwarded, or composed emails are transferred to your outbox the next time you synch. SD and CF slots allow upgrades and addition of software.

Email all your sysadmin questions and problems to Hans Huberland: sysadminqa@rackspace.co.uk Every published query on these system admin pages wins a Rackspace T-shirt.

It's only 74mm wide by 138mm long and a svelte 18mm thick – so the Sharp Zaurus really is Linux in your pocket!



★ Star Question – PDA winner!

This issue's lucky winner is **Silicon Cowboy** – your new Zaurus PDA will be with you shortly!

Users need ssh solutions

Q I run a Linux web server and all my users want to either use FTP or Frontpage to upload their web content. This is less than ideal as FTP uses clear text passwords, and Frontpage has security issues and crashes all the time. I've heard of sftp but there seem to be a bunch of different client versions or protocols out there. What should I do for a solid dependable client file transfer solution that just works – especially considering that most of my customers need drag-and-drop functionality and are running Windows home/client PCs?

Silicon Cowboy, via email

A Great points, and good question. If you're running a Linux web server and just "want it to work", then *ssh* (Secure Shell)-based solutions are probably a best over all solution given today's default Linux installs and available client side tools. With the stock installed *OpenSSH*, you get built in *sftp* (*ssh*-based file transfer) as well as *scp* (*ssh*-based "copy") options. *sftp* makes for a good solution but there is often much confusion regarding client compatibility given the various *sftp* server/client protocols that are out there, as well as OS-dependent client side requirements.

OpenSSH's implementation *scp* on the server side seems to be the best solution at this time, and most Linux servers now come with *OpenSSH/scp* installed and ready to run. One big caveat up to now has been that running *scp* requires the user to have shell login privileges (eg */bin/bash*) and it's a good idea not to give this out in a web server environment! However, this requirement is no longer true. The general login shell support that *scp* normally requires (usually *bash* on Linux) can be 'replaced' by a drop-in binary called *scponly* (see www.scponly.com) or search <http://rpmfind.net> for the latest binary rpm). *scponly* allows a web server admin to grant home-dir or chrooted home-dir based file transfers that are both authenticated and

fully encrypted – without having to give the client actual shell login privileges. If you want to add *scponly* to your system as a valid shell replacement however, just be sure that after installing it to also attach an entry for it to the last line of */etc/shells* to validate it for system use... And of course replace the *scp* user's */bin/bash* entry in the */etc/passwd* file with */bin/scponly* like this:

```
...
lxf:x:500:500::/home/lxf:/bin/bash
user1:x:501:501::/home/user1:/bin/scponly
...
```

The way *scp* works is just like the *cp* command, except it's over the network and encrypted. A client side command line *SSH/SCP* file transfer session looks something like this:

```
# scp -rp root@oldweb.mydomain.com:/var/www /var
root@oldweb.mydomain.com's password:
bind.html 100% |*****| 5680 00:00
LICENSE 100% |*****| 2827 00:00
bind.html.ja.jis 100% |*****| 6340 00:00
bind.html.en 100% |*****| 5564 00:00
bind.html.fr 100% |*****| 6663 00:00
cgi_path.html.en 100% |*****| 4452 00:00
cgi_path.html 100% |*****| 4568 00:00
cgi_path.html.ja.jis 100% |*****| 4804 00:00
cgi_path.html.fr 100% |*****| 6120 00:00
```

Notice the *-r* switch that allows recursive content transfers, and the *p* gets UID/GID and permission settings! But don't worry, your not going to have to force your Windows users to use a scary-looking command line tool. Check out <http://winscp.vse.cz/eng> for more information.

switches to the RPM command:

--aid – This will automatically install packages which are required if you have them available. You should only really use this with the **-U** (upgrade) switch, which will upgrade if the

package exists or install if it does not. Here's an example:

```
rpm -U lynx-2.8.5-11-i386.rpm
```

will fail because perl-CGI is not installed, while

```
rpm -U --aid lynx-2.8.5-11-i386.rpm
```

will automatically install this dependency too.

--redhatprovides – Using this option with the **-q** switch causes the rpm database of all packages provided by Red Hat for this release to be queried.

For example:

```
rpm -q --redhatprovides
libcanna.so.1.0
should return:
Canna-libs-3.6.12
```

We hope that this helps you.

Answers

Right right-click

Q I recently installed Red Hat 8.0 on a small computer I had lying around. I am currently experiencing a major difficulty with my GNOME menu. I cannot figure out how to add additional program launchers. The documentation I have says an option is available as a right-click on the menu. My right-click menu does have most of the options as listed in my help files, but not the specific menu option to create new launchers.


A second method is to use the 'Get Started' launcher; this will allow me to walk the GNOME menu in a Nautilus window as a series of folders and files, however, if I attempt to add a launcher this way, I can enter the information for the launcher, program path, icon, etc. But, it will not actually add the launcher to the menu! This seems to be a Nautilus issue, as I can add launchers to the desktop, and to the KDE menu (in KDE only...) I have tried this as both a user and under root, neither will work. Could it be a file permissions issue? Where is the information stored for the GNOME menu (default location)? What should the permissions be to allow what I want to do as both user and root? I previously had a computer with Red Hat 7.0 without this problem, I even had custom menu items for specific users and root.

Cecil P. via email

A GNOME menu items are added by creating a 'app.desktop' file in /usr/share/applications/. The format of a .desktop file is as follows:

```
[Desktop Entry]
Encoding=UTF-8
Name=Calculator
Comment=Perform simple calculations
Exec=gnome-calculator
Icon=gnome-calc3.png
Terminal=false
Type=Application
StartupNotify=true
Categories=GNOME;Application;Utility;
X-GNOME-DocPath=gnome-calculator/gnome-calculator.xml
X-GNOME-Bugzilla-Bugzilla=GNOME
X-GNOME-Bugzilla-Product=gnome-utils
X-GNOME-Bugzilla-Component=gcalc
```

By copying one of the defaults and modifying it as appropriate, it is fairly



The GNOME menu allows custom launchers to be added to the menu for each user, or for all users.

straightforward to add extra menu items to the GNOME menu. This, of course, will need to be performed as the root user.

GNOME allows a user-specific menu to be created, although you should be able to do this via the GNOME panel menu without having to use Nautilus. If you continue to have problems, upgrading your GNOME packages using the Red Hat errata updates at redhat.com may solve the problems. Red Hat 8.0 was a very short-lived release of the Red Hat Linux distribution, so unfortunately many of the problems in it have been ignored now that RH 9.0 is available.

Scanner stuck

Q I'm running MDK 9.0 and I am having some trouble getting my Canon N650U USB Scanner to work. I've looked on the mostang website and there they state that the Plustek backend can be used with this scanner.

Following their instructions I've added the following to plustek.conf:

```
[usb] 0x04a9 0x2206
dev/usbscanner
```

I have also added the following to /etc/modules.conf, as shown in the VueScan article (LXF35):

```
options scanner vendor = 0x04a9
product = 0x2206
```

After making these changes, running

```
cat /proc/bus/usb/devices; gives:
C:* #lfs = 1 Cfg# = 1 Atr = 80
```

```
MxPwr = 500mA
```

```
I: If# = 0 Alt = 0 #EPs = 3 Cls =
ff(vend.) Sub = 01 Prot = ff Driver =
(none)
```

```
E: Ad = 81(I) Atr = 03(Int) MxPS = 4
lvl = 5ms
```

```
E:Ad = 82(I) Atr = 02(Bulk) MxPS =
64 lvl = 0ms
```

```
E:Ad = 03(0) Atr = 02(Bulk) MxPS
= 64 lvl = 0ms
```

```
T: Bus=01 Lev = 01 Prnt = 01 Port
= 01 Cnt = 02 Dev# = 3 Spd = 12
MxCh = 0
```

```
D: Ver =1.00 Cls = 00(>ifc) Sub =
00 Prot = 00 MxPS = 8 #Cfgs = 1
```

```
P: Vendor = 04a9 ProdID = 2206
Rev = 1.00
```

```
S: Manufacturer = Canon
```

```
S: Product = CanoScan
```

Am I right to think that the kernel module still hasn't loaded? If so, what else do I have to do to make it load?

John Knops, Australia

A None of the above would suggest that the module is not loaded, although nothing

suggests that it is. You can load it using *modprobe scanner*, which should output some useful information when you execute *dmesg*. If you install the 'hotplug' daemon, then it can load the scanner module for you when the device is plugged in. Alternatively one can have the *scanner* module loaded at boot time by having a *modprobe scanner* command in a rc.d script and the module will be loaded ready for when you plug in the device.

Scanner episode II

Q I was wondering if any other readers share my predicament (being a Linux-newbie who hasn't a clue!)

I am having major difficulty in configuring my scanner to run on Mandrake 9.1 (from your cover DVD). It is one of those Packard Bell USB offerings that is really a Mustek Bearpaw 1200 CU Plus and is correctly identified by HardDrake. My webcam on the other hand, is listed as an OV511 device under scanners in HardDrake. When I run *X-Sane*, a dialogue appears asking me to choose what the device is:

1 *Noname OV511 USB Camera*

virtual device [v411:/dev/video0] (which is my Creative Webcam3 USB); or

2 *Mustek BearPaw1200 CU Flatbed Scanner [gt68xx:/dev/usb/scanner0]*.

I select the latter since the former ain't a scanner and *X-Sane* spurts out the reply that it failed to open device

```
gt68xx:/dev/usb/scanner0:invalid
argument
```

Hmmm, I don't remember any mention of 'gt68xx'...

Am I missing something REALLY obvious, or is something else about that prevents me from doing all those mundane scanning chores that people ask you to do once they discover that you have one! Any help from LXF would be appreciated in order that I cure this HP Pavilion.453.uk from the virus called 'XP' that is currently residing on one of its hard disks. Thanks in advance,

Dom Dorris, Glasgow

A gt68xx is the kernel module which provides support for this device. The 'invalid argument' is quite a strange one, although without any error debugging output from *dmesg* or *syslog*, it's fairly difficult to establish specifically what the error is caused by. It could be a problem with the kernel module, with the USB subsystem modules on the system, or any number of other problems.

If you use the '*scanimage*' tool with *scanimage -L* it will produce some fairly verbose output which may assist you in diagnosing the issue. If you manage to obtain some debugging information from the kernel, send it into LXF Answers and we'll take another look. [LXF](http://www.linuxformat.co.uk)

missed one?

LINUX FORMAT BACK ISSUES

Every month *Linux Format* brings you the best tutorials, the essential reviews and the latest news. But if you've missed out on a must-read feature or a vital program from our expertly compiled CDs and DVDs, order your back issue NOW! And remember, you need never miss an issue of your favourite Linux mag, if you subscribe to *Linux Format* (see overleaf for more details).



August 2003

Product code:
LXFB0043(cd)
LXFD0043(dvd)

DVD HIGHLIGHTS:
GnuCash, 3dDesktop, Slackware Live CD, LDP, Astaro Security Linux, SharpROM v3, Cinelerra, ScummVM, TrustixSL

MAGAZINE FEATURING:
Networking feature, VMware review, CD-burner roundup, Evesham ELinux desktop PC, Opteron servers, Linux training, TextMaker 2002, NetVault 7, GnuCash, 8 beginners' pages

CDs HIGHLIGHTS:
Complete Networking Toolkit, SpamAssassin, AGT, GuardDog, Jwall, Nagios, America's Army, Samba (stable & dev)



July 2003

Product code:
LXFB0042(cd)
LXFD0042(dvd)

DVD HIGHLIGHTS:
Opera 7.1 beta, GAIM, ABCDE, Pingus, KDE-Sounds/Aqua Fusion, LDP, Alsaplayer, KOffice, RUTE, Seclude, Freevix, DamnSmallLinux

MAGAZINE FEATURING:
All about X, MySQL review, Sharq server review, WINE explained, Steganography, Distributed Computing, Red Hat Enterprise AS, Choosing a distro, GCC 3.3 review

CDs HIGHLIGHTS:
Morphix, GuardDog, NVIDIA XDrivers, GCC 3.3, TORCS, Linux Cookbook, SETI



June 2003

Product code:
LXFB0041(cd)
LXFD0041(dvd)

DVD HIGHLIGHTS:
Slackware 9.0, Mutt, OpenOffice.org, Xine, TakeAJoint, Peanut Linux, FCKeditor, KDE 3.1.1a, mirror of Linux Documentation Project

MAGAZINE FEATURING:
AMD's 64-bit Opteron CPU, Linux Laptop Megatest, The Hurd, JCodeBox review, Build your own RPMs, Red Hat 9 vs SuSE8.2 vs Mandrake 9.1, OpenAuth Linux PC review

CDs HIGHLIGHTS:
Mandrake 9.1, AC3D, KDERadio, The GIMP - including 400 plugins from The GIMP Registry



May 2003

Product code:
LXFB0040(cd)
LXFD0040(dvd)

DVD HIGHLIGHTS:
Vector Linux, Dropline GNOME, Lopster, jEdit, QtParted, QuickRipDVD, BaitAndSwitchHoneyPot, AdiosBootCD, zTappy, MondoRescue

MAGAZINE FEATURING:
SCO v IBM \$1 billion lawsuit, Blender tutorial, LXF Awards 2002 results, Linux on Apple Mac, Gigabyte SR113 server, EFI to replace BIOS, VariCAD & IEMS mailserver reviews

CDs HIGHLIGHTS:
Mozilla 1.3, Opera update, IMP server, Kino, IceWM, ECLiPt Roaster, VariCAD



April 2003

Product code:
LXFB0039(cd)
LXFD0039(dvd)

DVD HIGHLIGHTS:
TensorNet, K3b, VIPS, FrozenBubble, Linux Open Media Box, IDMS Linux, WhiteDune, OBM, AstaroSecurityLinux

MAGAZINE FEATURING:
Customising Linux, Evesham Linux PCs, Open publishing with WIKI, Linux PIMs and PDAs, SME Server review, File Managers Roundup, Serial Networking, GnuCash

CDs HIGHLIGHTS:
SME Server, TightVNC, GNOME 2.2, VOCP, CoyoteLinux, UT2003 update, Blender, Qt, BootEverywhereLinux



March 2003

Product code:
LXFB0038(cd)
LXFD0038(dvd)

DVD HIGHLIGHTS:
Python, The GIMP, Lycoris, Yellowdog, Knoppix, Cygwin, DansGuardian, Senken, AbiWord, X Northern Captain, Eclipse

MAGAZINE FEATURING:
You, Copyright & The Law, X apps on Mac OS with Fink, SVG versus Flash, Audio Editors roundup, Python tutorial, GIMP secrets, SuSE Office Desktop examined

CDs HIGHLIGHTS:
KDE 3.1, UNIX Tools For Windows, Konstruct, phpMyAdmin, X Net Strength, AutoUpdate

Other back issues also available, please call for further information, giving the issue number and/or date required

To order, call the back issues hotline on **0870 4448645**

or **+44 1458 271178** for overseas orders

Remember to quote the issue code shown above and have your credit card/switch card details ready.

Prices (inc post and packing): CD – UK£5.99, US & Eur £6.99, rest of world £7.99
DVD – UK £6.99, US&Eur £7.99, rest of world £8.99.

LINUX

www.linuxformat.co.uk **FORMAT**

SUBSCRIBE TODAY AND RECEIVE YOUR FREE **LINUX** **HANDBOOK** WORTH £14.99!

HURRY – LIMITED STOCK

CALL NOW ON 0870 444 8645

ORDER FORM:

Yes! Please send me 13 issues of *Linux Format* magazine and my FREE *Linux Handbook*.
I understand this offer is only available if I subscribe by Direct Debit and am a resident in the UK.

Title _____ Initials _____ Surname _____
Address _____
Postcode _____ Country _____
Email address _____
Daytime telephone number _____



Media with passion
Future Publishing Ltd, Cary Court,
Somerton, Somerset, TA11 6BR



Instruction to your Bank or Building Society to pay Direct Debits

Originator's Identification Number

7 6 8 1 9 5

Please send this form to address shown

1. Name and full postal address of your Bank or Building Society branch

To: The Manager

Bank/Building Society

Address

Postcode

2. Name(s) of account holder(s)

3. Branch sort code

(from the top right-hand corner of your cheque)

4. Bank or Building Society account number

5. Instruction to your Bank or Building Society

Please pay Future Publishing Direct Debits from the account detailed on this Instruction subject to the safeguards assured by the Direct Debit Guarantee. I understand that this instruction may remain with Future Publishing and if so, details will be passed electronically to my bank or building society.

Signature(s)

Date

Ref No (Office use only)

Banks and Building Societies may not accept Direct Debit Instructions for some types of account.

You may photocopy this blank form and then complete it. You cannot fax or email Direct Debit forms.

QUOTE ORDER No. **LXFP44**

LINUX

FORMAT

THE UK'S BEST-
SELLING LINUX
MAGAZINE

Please choose your method of payment **1** or **2**

- 1.** ☐ Direct Debit (UK only – double CD Edition) – I wish to pay £13.73 every 3 months
or ☐ Direct Debit (UK only – DVD Edition) – I wish to pay £14.98 every 3 months

Please complete the form on the left

I understand that I will receive 13 issues during the next 12 months.

2. Cheque/credit card 13 issues

	CD	DVD
UK	<input type="checkbox"/> £54.90	<input type="checkbox"/> £59.90
Europe (including Eire)	<input type="checkbox"/> £78.90	<input type="checkbox"/> £83.90
North America	<input type="checkbox"/> £78.90	<input type="checkbox"/> £83.90
Rest of the World	<input type="checkbox"/> £91.90	<input type="checkbox"/> £96.90

☐ Cheque (payable to Linux Format. Sterling cheques drawn on a UK account.)

☐ Visa

☐ MasterCard

☐ American Express

☐ Switch

Issue No/Valid date

Card No. _____ Expires _____

Signature _____ Date _____

UK READERS RETURN THIS COUPON TO:

Linux Format, Subscriptions, Future Publishing Ltd, Freepost BS4900, Somerton, Somerset, TA11 6BR.

OVERSEAS READERS RETURN (POSTAGE PAYABLE) TO:


Linux Format, Subscriptions, Future Publishing Ltd, Cary Court, Somerton, Somerset, TA11 6TB, UK

WEB: www.linuxformat.co.uk then click on **Subscribe**

☐ Please tick this box if you would prefer not to receive information on other offers.

ORDER HOTLINE ☎ 0870 444 8645
OR FROM OVERSEAS ON +44 870 444 8645

OFFER ENDS: 08.09.03



YOUR FREE LINUX HANDBOOK:

- ✓ Essential 250-page guide to mastering Linux – whatever your level
- ✓ Everything you need to know about Linux – from installation to compiling code and setting up your own network
- ✓ Yours FREE when you subscribe today by Direct Debit*

**Handbook + 13
Linux Format
mags (DVD) =
SAVING OF
£39.46!**



(This offer is subject to availability and Future Publishing reserves the right to substitute this gift with an alternative)

- **THE ESSENTIAL** Linux read for all Linux users – can you afford to miss out?
- **EVERY ISSUE** comes with the ideal supplement for IT professionals – *Linux Pro* covers trends, technology, security issues, industry comment and thought-provoking case studies
- **GUARANTEED DELIVERY** direct to your door so you never miss an issue!

RISK-FREE GUARANTEE!

You can cancel your subscription at any time and receive a full refund on all unmailed issues, no questions asked

SPREAD THE COST OF SUBSCRIBING

For under £15.00* every 3 months you can pay by convenient Direct Debit. Simply complete the form opposite and return it back to us today.

Call now on 0870 444 8645 to make sure your receive the next issue Please quote LXF44

* Prices and savings based on DVD edition.

www.linuxformat.co.uk

LXF44 SEPTEMBER 2003 **103**



Essential disc info

Read this important information before you use your *Linux Format* coverdisc – CD or DVD.
We've collated some helpful info to help you get the most from these jewels of data!

FINDING THE ESSENTIALS

Missing something?

As many of the programs on our discs are the very latest releases, they are often built on the very latest libraries and may depend on other packages your current Linux setup does not contain. We try to provide you with as many of these important supporting files and libraries as possible, though obviously we don't have space to include absolutely everything.

In many cases, the latest libraries and

other packages you might need will be included in the "essentials" folder on the disc, so if you are missing dependencies, this is the first place to look.

Package formats

Wherever possible, we try to include as many different types of package for an installation as possible, whether that be distribution specific RPMs, debs or whatever. Please bear in mind that we can only do this where space permits and when the packages are available.

We will, apart from exceptional or legally restricted situations, include the source files for any package, so that you can build it yourself.

Documentation

These pages provide helpful information on how to install and use some of the packages on the CD. Please note that many of the applications come with their own documentation, and there are additional notes and files in the relevant directories.

WHAT ARE ALL THESE FILES?

If you are new to Linux, you may find the profusion of different files and extensions confusing. As we try to give as many packages as possible for compatibility, there will often be two or three files in a directory covering different types of Linux, different architectures and usually source and binary versions – so which do you install? They can be identified by their filenames, and usually just by the file extensions.

Someap-1.0.1.i386.rpm – This is probably a binary rpm, designed to run on x86 systems.

Someap-1.0.1.i386.deb – The same, but a debian package.

Someap-1.0.1.tar.gz – This is usually source code.

Someap-1.0.1.tgz – Same as the above, tgz is abbreviated form of tar.gz

Someap-1.0.1.tar.bz2 – Same, but uses bzip2 compression instead of zip

Someap-1.0.1.src.rpm – This is also source code, but supplied as an rpm to make it easier to install

Someap-1.0.1.i386.RH7.RPM – A binary, x86 RPM designed specifically for Red Hat Linux

Someap-1.0.1.ppc.Suse7.rpm – A binary RPM designed specifically for SuSE7.x PPC Linux.

Someap-devel-1.0.1.i386.rpm – A development version.

INSTALLING FROM TARBALLS

A tar ball is a two stage archive. First the files are archived into a single file with tar and then compressed with Gzip or Bzip2. To unpack, cd to the directory you want to unpack it, usually your home directory and type one of the following two lines:

```
tar xzvf /mnt/cdrom/Desktop/progname/progname-2.1.0.tgz
```

```
tar xvf -bzip2 /mnt/cdrom/Desktop/progname/progname-2.1.0.tar.bz2
```

Use the first for Gzipped files, those ending in .tar.gz or .tgz, and the second for Bzipped files, ending in .tar.bz2 or .tbz2. Naturally, you change the paths to suit the location and name of the archive. and replace /mnt/cdrom with whatever is applicable to your system (eg /cdrom). This normally unpacks the archive into a directory of the same name, enter that directory with:

```
cd progname-2.1.0
```

To compile and install the software, type the following three commands:

```
./configure
```

```
make
```

```
su -c "make install"
```

The last line will prompt you for the root password, as this stage must be run as root. If you are already logged in as root, just type **make install**. This will give you a default installation. If you want to change any aspect of the install, type

./configure --help to see the options available. For example, you are usually able to change the default location with the PREFIX argument. When you have finished installing, you may remove the source files with:


```
cd ..
```

```
rm -fr progname-2.1.0
```

You should also log out as root, before you do anything you may later regret.

DEFECTIVE CDs

In the unlikely event of your disc being defective please email our support team (support@futurenet.co.uk) for further assistance. If you would prefer to talk to a member of our reader support team please call **01225 822 743**.



CREATING INSTALL CDS WITH CDRECORD

The quickest way to burn an ISO image to CD is with *cdrecord*. You need to be root to do this. First find the address of your CD-writer with

```
cdrecord -scanbus
```

This will show the devices connected to your system. The SCSI address of each device is the three numbers in the leftmost column, say 0,3,0. Now you can burn a CD with

```
cdrecord dev=0,3,0 -v  
/path/to/image.iso
```

You can simplify the command by saving some default settings in /etc/default/cdrecord. Add a line for each CD writer on your system (usually one) like this

```
Plextor= 0,3,0 12 16M
```

The first item is a label, after the SCSI address you put the speed and the buffer size to use. You can now replace the SCSI address in the command line with the label, but it gets even easier if you add

```
CDR_DEVICE=Plextor
```

Now you can burn an ISO image to disc with

```
cdrecord -v/path/to/image.iso
```

If you really don't want to use the command line, *gcombust* will do the job for you. Start it as root, select the "Burn" tab and the "ISO 9660 Image" gadget near the top of the window. Put the path to the image file in the gadget and press "Combust!". Now put on the kettle while the CD is created for you.

Other OS?

You do not have to use Linux to burn the ISO to a disc. All the Linux-specific bits are already built into the image file. Programs like *cdrecord* simply dump it to the disk. If you don't have a CD-writer, find someone who does have one, and a DVD drive, and use the CD burning software on their computer. It can be Windows, Mac OS, AmigaOS whatever.

No CD burner?

What if you have no CD writer? Do you know someone else with one? You don't have to use Linux to burn the CDs, any operating system that can run a CD-writer will do the job (see above).

With some distributions it is also possible to mount the images and do a network install, or even a local install from another disk partition. The methods often vary between distributions, so check on the distro vendors website for more information. [LXF](#)

Coverdisc



Neil Bothwick is your guide through the wonders of this month's jam-packed *Linux Format DVD*. There's *even more* multimedia progs on here than the CDs!

We haven't included a proper Red Hat distribution on the cover discs for quite some time, because of the licensing and trademark issues involved. This month we have the next best thing, a distribution based on Red Hat. JAMD (Just Another Modified Distribution) is based on Red Hat and aimed specifically at the home desktop user. This means there are several differences between JAMD and Red Hat. There are no server or similarly business-oriented packages. It comes on a single CD, partly because of the omission of the server packages, partly because of reduced documentation. It includes multimedia software not included in a standard Red Hat

distribution, like MP3 and DVD playback programs. It is also easier to install, not because there is anything necessarily wrong with Red Hat's installer, but because there are so few choices that need to be made when installing a distribution intended for one particular group of users.

JAMD is optimised for i686 processors, this includes Pentium II and higher, AMD Athlon/Duron and Celeron CPUs. It does not include AMD K6 or VIA C3 processors, which fall somewhere between the older i586 specification and a full i686. If you try to install on an unsuitable machine, the installer will fail or even refuse to start.

JAMD LINUX INSTALLATION

The installation process for JAMD varies according to whether you have free space available on your hard disk. This means an area that is not partitioned, rather than free space on an existing partition. Despite what the introductory screen implies, it is not possible to resize existing partitions to make space from within the installer. You need to do this before running the installer. If you



No home-oriented distribution would be complete without a healthy selection of games, and JAMD Linux is no exception.

already have some form of Linux installed, you can use your distro's partitioning tool, or something like parted or qtparted. If you are currently Windows only, use something like *Partition Magic* to free up some space. Alternatively, you can boot from the JAMD CD, type **linux rescue** to get a basic shell and run parted. Press **h** while

parted is running to see the options and h followed by a command name for details on that command. JAMD needs around 2GB minimum for its default installation.

Now boot from the CD. Normally you would just press **Return** at the Boot prompt. There are various options that you can read about by pressing **F2** when the boot prompt is visible. After confirming some basic details about your system, you will be asked to provide a root password, this is needed for system administration, so don't forget it. Now the installer will prepare the empty space on your hard drive and begin installing the software packages. This can take a while depending on the speed of your computer; it took around three quarters of an hour on the 366MHz IBM laptop I use for testing, so you may want to put the kettle on now.

JAMD LINUX RUNNING FOR THE FIRST TIME

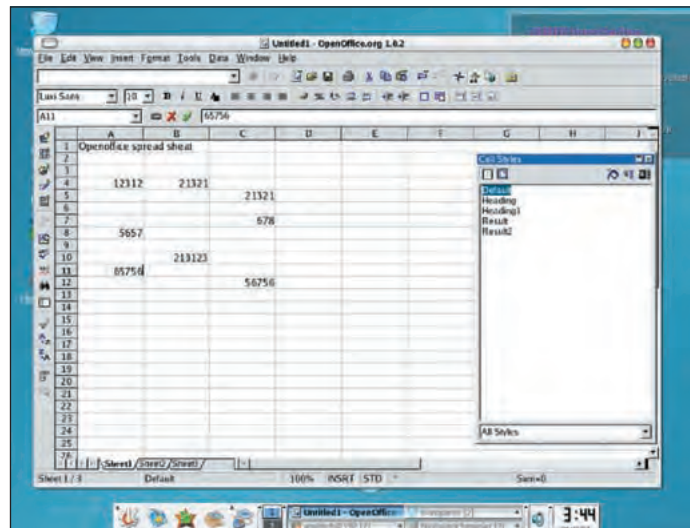
Once the installation is completed, remove the CD and let the computer reboot. The first time you use JAMD Linux you will be able to do a little



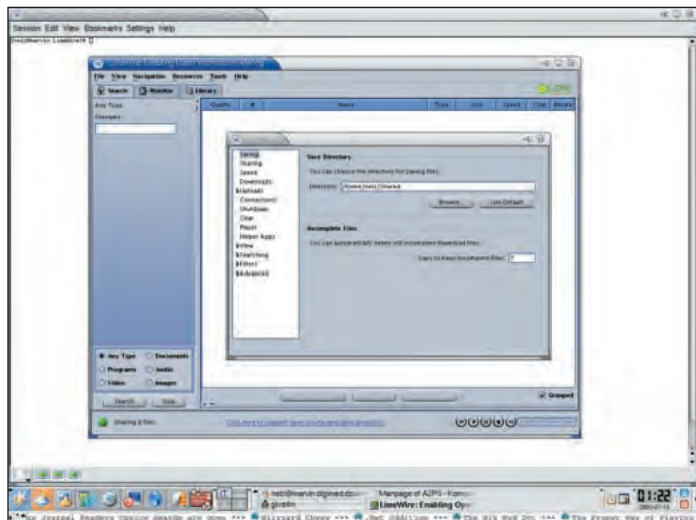
Wherever you see this logo it means there's related stuff on the DVD

IMPORTANT NOTICE

Before you even put the DVD in your drive, please make sure you read, understand and agree to the following: The *Linux Format DVD* is thoroughly tested for all known viruses, and is independently certified virus-free before duplication. We recommend that you always run a reliable and up-to-date virus-checker on ANY new software. While every care is taken in the selection, testing and installation of DVD software, Future Publishing can accept no responsibility for disruption and/or loss to your data or your computer system which may occur while using this disc, the programs or the data on it. You are strongly advised to have up-to-date, verified backups of all important files. Please read individual licences for usage terms.



JAMD Linux is not all play. It also includes software suitable for home productivity, like *OpenOffice.org*.



P2P file sharing with *LimeWire*, but you'll need Java installed first.

more configuring, this includes the most important step of creating a user account for yourself. The root account is reserved for system administration. In order to fit everything onto a single CD, there isn't a great deal of documentation, so take a look at www.jamd-linux.com/docs/ for extra information. There are also web forums at this site, where developers and other users are willing to help with your questions.

JAMD Linux comes with *Ximian Evolution*, the all-in-one email, PIM and calendaring program and you can play music files through *XMS*. If you use and enjoy JAMD, let the developers know. An email of acknowledgement would be nice, but what they really like is a postcard. That's not too much to pay for a complete operating system, is it?

**INTERNET
LIMEWIRE**

While there are legitimate concerns about some of the uses of point-to-point (P2P) file sharing, and a lot more FUD spread by those with an axe to grind, there is no doubt that P2P is here to stay, for now at least. There are several Linux P2P programs, but this is one of the better ones. *LimeWire* is a Java program, so you will need a Java runtime environment installed on your computer. Sun's own *J2RE* package or the *Blackdown* version are usually considered best, but we cannot include either of these on the DVD because of the Sun licence. If you are unable to download or otherwise obtain one of these, you can install *Kaffe* from the System directory of this DVD.

Provided you have Java installed, installing *LimeWire* is simply a matter of

opening a console and typing

```
sh /mnt/cdrom/Internet/LimeWire/  
LimeWireLinux.bin
```

The installer is designed to be run as an individual user, putting everything in your home directory. Even if you install as root, it just puts it all in root's home directory. Unless you particularly need to share one *LimeWire* installation between several users, in which case you should install it somewhere like `/usr/local/share` and add its directory to the path, it is best to install it as your normal user.

MULTIMEDIA DISTROS/MOVIX

The Multimedia directory contains all the contents of the special multimedia themed CD, plus a few goodies reserved for our DVD readers. One of these is the innovative Movix mini-distro. Movix is a distro designed to run only one program, *MPlayer*. The idea is that you put Movix and your favourite video files onto a CD-R or DVD-R and you can play the videos on just about any X86 computer, without installing anything to its hard drive.

There are three varieties of *MoviX*. *MoviX* itself boots from CD into a console, from where you may run *view* any of the files on the CD with *MPlayer*. *MoviX2* is similar to *MoviX*, but it includes *X* so you can use *MPlayer's* GUI. *eMoviX* is designed to be written to a CD together with all video/audio files you want, so that the CD will boot and automatically play all files. This makes *eMoviX* ideal for distributing videos as self-playing CDs to anyone with a PC, whatever operating system they use.

MoviX and MoviX2 are on the DVD as ISO images (compressed with Zip) as well as source code, so you can create a

Advanced mode install

More control – Personal Desktop or Custom

By typing "advanced" at the boot prompt, you will be given a lot more control, and responsibility, over the install process. You can decide exactly how your hard disk is partitioned. Advanced more gives a choice of two installations: Personal Desktop or Custom. The former installs the complete selection of packages for normal desktop and laptop usage, as you get with the standard installation. The latter lets you make your own choices instead. Whichever you go for, you can always add and remove

packages later, once your system is up and running.


The advanced installer allows you to override the default partitioning system with the *Disk Druid* tool, set a security level and decide which services to allow, which you should leave off as **JAMD** doesn't install servers. When the package installation is complete you can adjust the configuration of *X*, the basis of the graphical interface, to work with your video card. Then you can set up the bootloader and create a boot disk.

MoviX CD just by unzipping one of the ISOs and burning it to a CD-R. eMoviX is only available as source code, which you then use to build your own ISO image that includes your movie files. Full instructions are included in the tarball.

OFFICE
KOFFICE

We recently had a beta version of the forthcoming *OpenOffice.org* on the DVD, now it's the turn of the KDE guys for a share of the limelight. Development of *KOffice 1.3* is well advanced and a second beta version has been released. The DVD carries the source code for this as well as packages for Red Hat, Mandrake and SuSE. Because of space limitations, we have only included locale files for the

most common languages. If you want to use one of the many other languages available, either install from source (maybe using *Checkinstall* from the Essentials directory to create a package) or download the relevant package from kde.org. They are all fairly small downloads, but the combined size of all the languages available for each of the three distros was more than the DVD could bear.

If your needs are more for desktop publishing than word-processing, we have the release candidate for version 1.0 of the *Scribus* DTP package. 1.0 itself was released just too late to make it onto the DVD. With the two releases so close together, the release candidate on the DVD is likely to be very close to the finished package. 

Finding files on the CDs and DVDs

Using the Index

What happens when you want a particular program that you think was on a previous cover disc? You could try digging out all your discs, putting each one in the drive and trawling through the index files for what you need. Although you would eventually find it, it would probably have been quicker to download it (unless it's a 650MB ISO image). There is a much faster way, two in fact.

If you have a Java-enabled web browser, load the index.html file of the current DVD and you'll see a search box towards the top right of the page. Type in part of the program name and press return and it will search the recent discs. Select the program you want from the Results window, making sure the appropriate disc is mounted, and it will load the program's directory into the browser window. There are two check boxes below the search gadget, the first restricts the search to the current disc only and is selected by default. The other causes the applet to search the descriptions of each program as well as the names. This is useful when you aren't sure of the name of the program and want to search on what it does.

The searching is fairly simple at the moment, a straight substring search. If you enter more than one word, it will search for that phrase, rather than for descriptions containing any combination of those words. This is a fairly new addition to the cover discs, so the index files for these only go back a few months. The alternative, which has been available since issue 20, is a set of CSV files in the Essentials/CSV directory of the disc. These are ready for import into any database program, with each line of the file containing the following data:

Disc.	Name	Path	Description	Home page	Version
1
2
3
4
5
6
7
8
9
10
11
12
13
14
15
16
17
18
19
20
21
22
23
24
25
26
27
28
29
30
31
32
33
34
35
36
37
38
39
40
41
42
43
44
45
46
47
48
49
50
51
52
53
54
55
56
57
58
59
60
61
62
63
64
65
66
67
68
69
70
71
72
73
74
75
76
77
78
79
80
81
82
83
84
85
86
87
88
89
90	

Or, if you're really lazy, there's another option. You can simply search the CSV files directly with **grep**, it works for me.

User Groups

LUGs worldwide are full of members keen to help with your problems, discuss ideas, and generally natter about all things Linux. You can find lots more information online at: www.lug.org.uk

1 HAMPSHIRE

URL www.hants.lug.org.uk
Contact Hugo Mills

2 BRISTOL & BATH

URL www.bristol.lug.org.uk

3 SCOTTISH

URL www.scottish.lug.org.uk

4 OXFORD

URL www.oxford.lug.org.uk
Contact Alasdair G Kergon

5 KENT

URL www.kent.lug.org.uk
Contact Kevin Groves

6 BRIGHTON

URL www.brighton.lug.org.uk
Contact Johnathan Swan

7 WORCESTERSHIRE

URL www.worcs.lug.org.uk

8 NORTHANTS

URL www.northants.lug.org.uk
Contact Kevin Taylor

9 ANGLIAN

URL www.anglian.lug.org.uk
Contact Martyn Drake

10 MILTON KEYNES

URL www.mk.lug.org.uk
Contact Denny De La Haye

11 SCUNTHORPE & DONCASTER

URL www.scundog.org
Contact Shaun Holt – shaun@scundog.org

REVISED
DETAILS

12 MORAY

URL www.moray.lug.org.uk
Contact Stewart Watson

13 WEST WALES

URL www.westwales.lug.org.uk
Contact Dan Field

14 WOLVES

URL www.wolveslug.org.uk
Contact Jono Bacon

15 PETERBOROUGH

URL www.peterboro.lug.org.uk
Contact Steve Gallagher

16 EDINBURGH

URL www.edinburgh.lug.org.uk
Contact Alistair Murray

17 TYNESIDE

URL www.tyneside.lug.org.uk
Contact Brian Ronald

18 LEICESTER

URL www.leicester.lug.org.uk
Contact Clive Jones

19 GREATER LONDON

URL <http://glug.linux.co.uk/>
Contact John Southern

20 SURREY

URL www.surrey.lug.org.uk
Contact Jay Bennie

21 CAMBRIDGE

URL www.cam-lug.org.uk

22 DEVON & CORNWALL

URL www.dclug.org.uk
Contact Simon Waters

23 FALKIRK

URL www.falkirk.lug.org.uk

24 MANCHESTER

URL www.manlug.mcc.ac.uk
Contact John Heaton, Owen Le Blanc

25 HERTFORDSHIRE

URL www.herts.lug.org.uk
Contact Nicolas Pike

26 WEST YORKSHIRE

URL www.wylug.lug.org.uk
Contact Jim Jackson

27 SHEFFIELD

URL www.shefflug.co.uk
Contact Richard Ibbotson

28 STAFFORDSHIRE

URL www.staffslug.org.uk

29 NORTH EAST

URL www.shofar.uklinux.net/NELUG

30 LONDON

URL www.lonix.org.uk

31 BERKSHIRE & THAMES VALLEY

URL www.sclug.org.uk

32 LIVERPOOL OPENSOURCE

URL http://linux.liv.ac.uk/_liv_linux_ug/
Contact Simon Hood

33 DEAL AMIGA CLUB

Email superhighwayman@hotmail.com
Contact John Worthington

34 CHESTERFIELD

Email spirelug@yahoo.co.uk
Contact Robin Needham

35 SOUTH DERBYSHIRE

URL www.sderbylug.org.uk
Contact Dominic Knight

36 BELFAST (BLUG)

URL www.belfastlinux.cx
Email russell@belfastlinux.org

37 WILTSHIRE

URL www.wiltshire.lug.org.uk
Contact Jason Rudgard

38 SOUTH LONDON

URL www.sl.lug.org.uk
Email edo@perceptiondm.com

39 CHESHIRE

URL www.sc.lug.org.uk
Contact Anthony Prime – enquiry@sc.lug.org.uk

40 NORTH WALES

URL www.northwales.lug.org.uk
Contact Andy Hutchings A-Wing deltaone@virgin.net

41 MIDLANDS

URL <http://midlandslug.port5.com/>
Contact Pete Thompson

42 CUMBRIA

URL www.cumbria.lug.org.uk
Contact Jamie Dainton

43 DORSET

URL www.dorset.lug.org.uk
Contact John Robinson

44 SHROPSHIRE

URL www.shropshire.lug.org.uk
Email shropshire@lug.org.uk

45 SOUTH WEST

URL www.southwest.lug.org.uk
Email southwest@lug.org.uk

46 SOUTH WALES

URL www.swlug.org.uk

47 NORTH LONDON

URL www.kemputing.net/lug/anlug-aims.html
Email jason@voyagercomputers.co.uk

48 MALVERN

URL www.malvern.lug.org.uk
Contact Greg Wright

49 HUDDERSFIELD

URL www.hud.lug.org.uk
Contact Dave Naylor – knocker@caramboo.com

50 NOTTINGHAM

URL www.nottingham.lug.org.uk

51 ST ALBANS & LUTON

URL www.lust.lug.org.uk
Contact Michael Culverhouse – mike@easily.co.uk

52 WREXHAM

Contact Paul Kersey-Smith
Email paul@pkls.fsnet.co.uk

53 PRESTON & LANCS

URL www.preston.lug.org.uk
Contact Phil Robinson

54 DERRY

URL www.derry.lug.org.uk

55 ISLE OF WIGHT

URL www.iow.lug.org.uk
Contact David Groom – info@iow.lug.org.uk

56 SCARBOROUGH

URL www.scarborough.lug.org.uk

57 BLACKBURN

Email matt@consultmatt.co.uk

58 YORK

URL www.york.lug.org.uk

59 LINCS

URL www.lincs.lug.org.uk

**60 HULL**URL www.hull.lug.org.uk**61 WALTON-ON-THAMES**Contact William Mutch
Email rael@freeuk.com**62 GLOUCS & COTSWOLDS**URL www.gloucs.lug.org.uk**63 WEST OF SCOTLAND**URL www.wos.lug.org.uk**64 SOUTH STAFFORDSHIRE**URL www.staffs.lug.org.uk**65 MANSFIELD**URL www.mansfield.lug.org.uk**66 BORDERS**URL www.linux.bordernet.co.uk**67 BIRMINGHAM**URL www.sb.lug.org.uk**68 COVENTRY**Email info@coventry.lug.org.uk**69 NEWARK & LINCOLN**URL www.newlinc.lug.org.uk**70 BEDFORDSHIRE**URL www.beds.lug.org.uk**71 LINCOLN**URL www.lincoln.lug.org.uk**72 LOUGHBOROUGH**URL www.loughborough.lug.org.uk**73 EXETER UNIVERSITY**URL www.euslug.lug.org.ukEmail N.J.Murison@exeter.ac.uk**74 SUNDERLAND**Email thomas.croucher@sunderland.ac.uk**75 EAST YORKSHIRE**Email sharkonline@whatemail.com**76 CLEVELAND OPEN SOURCE GROUP**Email openlug@digitalmedia.co.uk**77 BEVERLEY**Email vladimir_lukyanov@hotmail.com**78 DUNDEE & TAYSIDE**URL www.dundee.lug.org.uk**79 SUSSEX**URL <http://sussex.lug.org.uk/>**80 WIGAN & ST HELENS**Email paulf.johnson@ukonline.co.uk**81 BRIXTON**URL www.communitytechnology.org.uk/~linuxhome**82 ST.ANDREWS, FIFE**URL www.standrews.lug.org.ukEmail stuart@nx14.com**83 NUNEATON**URL www.nuneaton.lug.org.uk**84 ISLE OF MAN**URL www.iom.lug.org.ukEmail helix@manx.net**85 AYLESBURY**URL www.aylesbury.lug.org.ukEmail drbond@educational-computing.co.uk**86 LANCASHIRE**URL www.lancasterlug.org.uk**87 EAST LONDON**URL www.eastlondon.lug.org.uk

Contact Jonathan Spriggs

88 ORMSKIRKEmail rob@northwestlinux.co.uk**89 HEREFORD**URL www.hereford.lug.org.uk/Email rbjh@good-news.fsnet.co.uk**90 EAST HERTS**Email madtom1999@yahoo.com**91 SWINDON**Email nick.trueman@ntlworld.com**92 MENAI**URL www.menai.lug.org.uk**93 ABERDEEN**URL www.aberdeen.lug.org.uk**94 SHETLAND**URL www.shetland.lug.org.ukEmail c_s_s_butler@yahoo.com**95 GLASTONBURY**URL www.glastonbury.lug.org.uk

Contact Steve Leonard-Clarke

96 SOUTHEAST-ON-SEAURL www.sos.lug.org.uk

Contact Derek Shaw

97 ORPINGTONURL www.orpington.lug.org.uk

Contact Barry Schofield

YOUNG LINUXURL www.young.lug.org.uk**SCHOOLS**URL www.schools.lug.org.uk

LUGS OF THE MONTH

Wireless LUG (Orpington)

Barry Schofield writes:

Orpington LUG has been wired for wireless for a while now, and many members took advantage of the free broadband connection to show many interesting websites and Linux goodies at our meeting on 1st July. Most found it easy to get connections, especially with the USB wireless units floating around.

Naturally, there was a great interest in extending the wireless range and homemade antennae was

a much-discussed topic. Some members, having seeing sites with info on how to make antennae from crisp packets voted to get right on to it! Sites we consider vital to wireless Linux are added to our site pointer list at www.orpington.lug.org.uk.

Prices on the wireless PCMCIA cards for the 22Mbps 802.11b+ are about £20 now, though the cards will not be tested until our next meet on 29th July 2003. Hopefully they will be as simple to install as the USBs

which worked with Mandrake 9.1 and various flavours of Red Hat on members' laptops. Maybe we'll have a SUSE laptop turning up at our next meet?

A star of the meeting was the DIY Linux PC for under £200 with a VIA board.

If you need more info or want to exchange tips about wireless Linux, please send emails to barry@orpington.lug.org.uk or phone 07970 100313.



Worldwide Linux User Groups

Free Software users across the globe

Africa

EGYPT

URL www.linux-egypt.org

GAUTENG, SOUTH AFRICA

URL www.glug.org.zaEmail glugmin@revolution.org.za

THE LORD'S ABODE, JO'BURG, SA

Email Andrew.Gargan@avrin17@iname.com

Australia

ADELAIDE

URL www.linuxsa.org.auEmail mtippet@anu.edu.au

ALICE SPRINGS

URL www.aslug.org.au

MELBOURNE, VICTORIA

URL www.luv.asn.auContact luv-committee@luv.asn.au

PERTH

URL <http://plug.linux.org.au/>

SYDNEY

URL www.slug.org.au

Europe

COSTA DEL SOL (English speaking)

URL www.fuengirola.lug.org.uk

DENMARK

Alssund www.alslug.dkEsbjerg www.eslug.dkFyns www.flug.dkMidt-og Vestjylland www.mvjlug.dkNordjylland www.njlug.dkSkåne Sjælland www.sslug.dkTrekantsområdet www.tlug.dkVest-fyn www.haarby-net.dk/vflugÅrhus www.aalug.dk

EIRE

URL www.linux.ieEmail root@linux.ieURL www.dilu.orgContact glossary@dilu.org

MILUG (Longford)

URL <http://midlands.linux.ie>Contact midlands@linux.ie

Middle East

ISRAEL

URL www.iglu.org.il/IGLU/Contact webmaster@iglu.org.il

PALESTINE

URL www.lugps.orgEmail isam@planet.edu

Asia

HONG KONG (multilingual)

URL www.linux.org.hk

SINGAPORE – SLUG

URL www.lugs.org.sg

SRI LANKA

URL www.lklug.pdn.ac.lk

MYANMAR (formerly BURMA)

URL www.myanmarlug.orgEmail aftyde@balug.org

PAKISTAN

URL www.linuxpakistan.netEmail tux@clug.org

HYDERABAD, SINDH, INDUS VALLEY

URL www.geocities.com/slug_pk/

KASHMIR

Coming soon!

China

BEIJING (GB encoding, but mostly written in Chinese)

URL <http://mud.263.net.cn/~linux>

CHINESE LINUX USER GROUP

URL www.linux.org.cn

NANJING

URL <http://jllib.jlonline.com/njlug>

India

LINUX INDIA

URL <http://linux-india.org>

ALIGARH LUG

URL <http://linux.amupost.com>

BOMBAY

URL www.ilug-bom.org.in

CHANDIGARH

URL www.geocities.com/vipinb

CHENNAI AND MADRAS

URL www.chennaiug.org/

CYBERABAD (CLUG)

URL <http://seeknew.freesevers.com/clug/>

DELHI

URL www.linux-delhi.org

KOLKATA

URL www.ilug-cal.org

MADURI

URL <http://linuxmadurai.tripod.com>

NORTHERN INDIA LINUX

URL <http://groups.yahoo.com/group/lug-northindia>

Spreading the word

Linux use in schools isn't only a good idea in terms of cost, it's our most important advocacy asset, says **Jono Bacon**.

I have received a few emails from school IT managers sharing more info about some new ideas and issues relating to Linux use in education.

It is important to remember that schools have a limited budget, and that money spent on licensing software is money that is not spend on stationary, books etc. It's also important to remember that many schools are given an IT or equipment budget, and if they don't spend it, they don't get it next time. As such, IT managers mostly don't risk the budget and keep with commercial licenses, although as one reader pointed out, this budget can often be spent on more tangible teaching technology like projectors, more PCs, media etc.

Most commercial software apps have Linux equivalents: *Microsoft Office* vs *OpenOffice.org*, *SQL Server* vs *MySQL*, *ASP* vs *PHP* etc. In the commercial sphere though, some servers require Client Access Licenses (CALs) to add users to this software. In this way, the addition of CALs means you pay more money to use the Software you have

already paid for! With Linux, there are no such things as CALs, and this point should be made clear as well as the alternative applications available.

Including the training requirements in your financial proposal is essential. A move from Microsoft Windows-based desktops/servers to Linux means re-training, and many IT managers worry about both being seen to 'rock the boat' by non-IT literate budget administrators, and being left in the lurch if trained staff then leave. When advocating Linux use, if you stress that the training can be paid for with the savings made in licensing, and that new staff members can quite easily be trained in-house, you are likely to be more successful.

Next month we will begin looking at charities as a target for free software. I have been doing a lot of work in this area recently, and together we can help save these vital organisations some money. Keep on advocating Linux, and let me know how you get on through linuxformat@futurenet.co.uk or www.jonobacon.org **LXF**

Linux User Group organisers

If you're not listed here, or we have your details wrong, please contact us at: **LUGS!, Linux Format, 30 Monmouth Street, Bath, BA1 2BW** or email your details to: linuxformat@futurenet.co.uk

EDITORIAL

Editor Nick Veitch nick.veitch@futurenet.co.uk
Reviews Editor Paul Hudson paul.hudson@futurenet.co.uk
Art Editor Julian Jefferson julian.jefferson@futurenet.co.uk
Production Editor Matt Nailon matthew.nailon@futurenet.co.uk

Editorial Contributors Jono Bacon, Mike Saunders, Chris Brown, Hoyt Duff, David Cartwright, Richard Cobbett, Richard Drummond, David Coulson, Hans Huberland, Andy Channelle, Michael J Hammel

COVER CD PRODUCTION

CD Editor Neil Bothwick

ART CONTRIBUTORS

Art assistance Martin Mulchinock
Photography Amanda Thomas, James Wilson, Corbis, Eyewire, Photodisc, Powerstock
Illustration Paul Bateman, Chris Winn, Shane Collinge

ADVERTISING SALES

Key account manager George Gill george.gill@futurenet.co.uk
Deputy portfolio Ad Manager Diane Clydesdale
Business development manager Damian Hughes
Senior Sales Executives Giles Crosthwaite-Scott, Harry Maltby
Classified Sales Executive Stephen Hall

MARKETING AND PROMOTIONS

Marketing Manager Mike Hawkins
Subscriptions Executive Marie Lock

PRODUCTION

Production Co-ordinator Diane Ross
Group Production Manager Clare Tovey
Advertising Production Co-ordinator Jo Crosby

MANAGEMENT

Publisher Kelley Corten
Group Publisher Dave Taylor
Publishing Director John Weir
Managing Director Colin Morrison

DISTRIBUTION AND CIRCULATION

Circulation Manager Jamie Malley
jamie.malley@futurenet.co.uk

Distributed by Seymour Distribution, 86 Newman Street, London W1T 3EX **Tel** +44 (0)207 3968000
Overseas Distribution by Future Publishing Ltd. **Tel** 01225 442244.

Overseas Licences

International Licensing Director Simon Wear
simon.wear@futurenet.co.uk **Tel** +44(0)1225 822798
Fax +44 (0)1225 788105 **Cell phone** +44(0)777 5641493

Contact Details

Linux Format, 30 Monmouth Street, Bath BA1 2BW
Tel +44 (0)1225 442244 **Email** linuxformat@futurenet.co.uk

Subscriptions and Mail Order

Phone +44 (0)1458 271178. See page 102
Email linuxformat.subs@futurenet.co.uk

Copyright No part of this publication may be reproduced without written permission from our publisher. We assume all letters sent – by email, fax or post – are for publication unless otherwise stated, and reserve the right to edit contributions. All contributions to Linux Format are submitted and accepted on the basis of non-exclusive worldwide licence to publish or license others to do so unless otherwise agreed in advance in writing. Linux Format recognises all copyrights in this issue. Where possible, we have acknowledged the copyright holder. Contact us if we haven't credited your copyright and we will always correct any oversight. We cannot be held responsible for mistakes or misprints.

All CD-ROM demos and reader submissions are supplied to us on the assumption they can be incorporated into a future covermounted CD-ROM, unless expressly stated to the contrary.

Disclaimer All tips in this magazine are used at your own risk. We accept no liability for any loss of data or damage to your computer, peripherals or software through the use of any tips or advice.

Printed in the UK by Midway Clark (Holt) and Southern Print
 © Future Publishing Ltd 2003

LINUX is a trademark of Linus Torvalds, GNU/Linux is abbreviated to Linux throughout for brevity. All other trademarks are the property of their respective owners

Future Publishing Ltd. is part of The Future Network plc.

The Future Network produces carefully targeted specialist magazines for people who share a passion. We aim to satisfy that passion by creating titles offering value for money, reliable information, smart buying advice and which are a pleasure to read. Today we publish more than 90 magazines in the UK, US, France and Italy. Over 80 international editions of our magazines are also published in 28 other countries across the world.

The Future Network plc is a public company quoted on the London Stock Exchange (symbol: FNET).

Non-executive Chairman Roger Parry
Chief Executive Greg Ingham

Chief Operating Officer & Managing Director, UK Colin Morrison
Group Finance Director John Bowman

Tel +44 (0)1225 442244

BATH • LONDON • MILAN •
 NEW YORK • PARIS •
 SAN FRANCISCO • WROCLAW
www.thefuturenetwork.plc.uk



NEXT MONTH

Issue 45 on sale Tuesday 9 September

FASTER FASTER...

In terms of computing, you can never really be motoring too fast. Next month our expert guide to tweaking your box will have you booting, starting and running applications faster than before, with detailed guides to startup, kernel, device, memory and software setup.

Server school

We promised, now we deliver – a new series whose ultimate goal is to tell you everything you need to know about running a Linux sever, from hardware configuration, to apps and services. Don't miss part one!

On test

Latest Opteron server from Armari, OpenGroupware.org, Yellow Dog Linux, Runtime Revolution, and more...

PLUS:

DON'T MISS NEXT MONTH'S LINUX PRO!

20 pages of real-world Linux news, views and case studies for IT professionals – next month: storage special!



The exact contents of future issues are subject to change

DON'T MISS YOUR COPY OF LINUX FORMAT!

Get it delivered to your door every month – subscribe on page 102

LINUXPRO

FROM THE MAKERS OF LINUX FORMAT

SEPTEMBER 2003



PLUS

EXCLUSIVE!

The new Zend Studio 3

"Debugging even the most complicated form is just one click away" Zeev Suraski, Zend CTO

Which Host?

Top dedicated server hosts vie for your custom

Compiler writing

Special Q&A session clears up compiler queries

Who's made off with your name?

Trademarks, the law and domain names under scrutiny – can you protect your online identity and who decides which names belong to whom? **p6**

PRACTICAL LINUX SOLUTIONS FOR I.T. PROFESSIONALS

Welcome

Twenty pages of real-world Linux for IT professionals

You may think that if you own a trademark, you'll automatically get the nod to have the associated domain names. Not so, just ask Armani. In fact, the rules governing the use and abuse of domain names, while not being completely vague, do leave a lot to the imagination of the panelists who make rulings in cases of disputes. So how can you make sure you get the name you want, and more importantly, how can you prevent the unscrupulous from hijacking your brand and trading off your name? Our major feature in this issue outlines the legalities surrounding domain names, with some tips for keeping hold of what you've got.

The Internet features again in this issue of *Linux Pro* in the form of our web hosting feature. You have probably seen the ads for various server hosting companies in the magazine. We wanted to try them out and give some idea of the range of services and support you can expect from the biggest names in Linux dedicated servers. No punches were pulled – we even went to the lengths of deliberately sabotaging the servers and then calling up tech support to see if they could fix them. Overall the results were impressive, so skip ahead to page 16 for the full lowdown if you use Linux on the Web.

We were planning to have an exclusive case study of how AMD uses Zend's PHP products to power their web presence. Sadly this fell through, but our loss is your gain, because in the end the nice people at Zend decided to let us have a full and exclusive interview on the future of *Zend Studio* and what we can expect in release 3.

Finally this issue, we have a Q&A special on our compiler writing series. From next month we'll be moving this to the *Linux Format* magazine, so if you want to continue following this entertaining series, make sure you read *LXF* as well as *Linux Pro*.

As ever, if you have any feedback on what you read hear, or would like to read, you can reach me at the address below.

Nick Veitch Editor
nick.veitch@futurenet.co.uk



“No punches were pulled – we even tried sabotaging the servers, then calling tech support to see how long it took them to spot the problem.”

Contacts

EDITORIAL

Editor
Nick Veitch nick.veitch@futurenet.co.uk
Reviews editor
Paul Hudson paul.hudson@futurenet.co.uk
Art editor
Julian Jefferson julian.jefferson@futurenet.co.uk
Production editor
Matt Nailon matthewnailon@futurenet.co.uk
Group publisher
Dave Taylor
Publisher
Kelley Corten
Contributors
David Harris
Photography
Getty, Powerstock

ADVERTISING

Key Account Manager
George Gill george.gill@futurenet.co.uk
Business Development Manager
Damian Hughes damian.hughes@futurenet.co.uk
Senior Sales Executives
Diane Clydesdale, Giles Crosthwaite-Scott, Harry Maltby
Classified Sales Executive
Stephen Hall

LinuxPro is a part of
LINUX
FORMAT
magazine

Published by
the future network
Media with passion

30 Monmouth Street, Bath BA1 2BW
Tel +44(0) 1225 442244

SEPTEMBER 2003

CONTENTS

ZEND STUDIO 3

The future of the best PHP development platform revealed exclusively **p4**

WHAT'S IN A NAME?

Domain name shenanigans – how can you protect your brand or stop a global global corporation hijacking your site **p6**

COMPILER WRITING

Q&A on the story so far **p14**

HOSTS HEAD-TO-HEAD

Which hosting service is best for your needs? **p16**



6

16

LINUXPRO 3

ZEND STUDIO 3



The future of PHP Editing

When you think “PHP”, “Zend” is never far away – since the launch of PHP 4 with the first version of the *Zend Engine* powering it, Zend have provided comprehensive support for PHP programmers and server admins to make sure they can optimise their PHP solutions, and now their products are virtually ubiquitous in the market.

Zend Studio (2.5 reviewed in *LXF31* and 2.6 in *LXF40* where it received an impressive 9/10 score) has long been the most popular PHP editor – not only is it the most powerful, but, thanks to its Java nature, it also works as smoothly on Windows as it does on Linux or Mac OS X. As such, when a new version of *Zend Studio* is announced, especially a major new version, it’s greeted with much anticipation.

Now that Zend is launching the first beta of *Zend Studio* 3, we spoke to Zeev Suraski, the CTO of Zend and one of the lead programmers of PHP itself, to see what all the fuss is about...

LXF: What are the big new features in Studio 3?

ZEEV SURASKI: It’s a major new version of the product, and there are several key features that we’ve added since 2.6. We now support PHP 5, at least to the level where it is right now. Hopefully it won’t change too much from now until release, as far as the language syntax is concerned. There’s also integrated profiling support, and static code analysis.

One of the coolest new features is integration of some of the Studio functionality into web browsers – into *Microsoft*

Not content with already being head and shoulders ahead of its competitors, Zend has been working hard to improve Zend Studio even more. PAUL HUDSON spoke to Zeev Suraski about Zend Studio 3...

Internet Explorer and hopefully *Mozilla* and *Netscape* too. This will enable people to significantly simplify the way debugging and profiling works.

LXF: Is it built on top of Studio 2.6, or is version 3 a ground-up rewrite?

ZS: It’s built on top of *Zend Studio* 2.6, taking advantage of the latest version of Java, which has some significant performance gains. There are also lots of coding optimisations in the *Studio* itself – one of the things that was completely rewritten was the code completion, as well as the whole back-end representation of the project. This had two side effects: one was that it was much faster, and the other, which was the goal really, was that it’s now much more robust. It now supplies more information for code completion – particularly for PHP 5.

LXF: What’s the biggest new feature in Zend Studio 3 from your point of view?

ZS: Well, you probably hear this a lot, but it depends. If it’s performance intensive, then there’s no doubt that profiling is going to be a very big feature – it’s very good profiling, trust me, not just a load of numbers. The profiler tries to help you out by giving a graphical representation of what’s going on, as well as the hard numbers.

Code analysis is also very interesting. It can find a lot of trivial and non-trivial mistakes in scripts, a basic example of which is that it can detect and warn about unused variables. I think we’ll see a lot of growth in this feature – detection for more and more bugs will be added.

LXF: And, of course, it handles PHP 5 smoothly...

ZS: Yeah, this is definitely a “must” for companies that want to upgrade to PHP 5, because it has full support for the new features. While you can edit PHP 5 code with *Studio* 2.6, it's not going to give you any particular insight into PHP 5-specific code. That's where *Studio* 3 wins – if you want PHP 5 support, this is clearly an important feature.

LXF: How advanced is the code analyser?

ZS: What we're talking about at this stage is a static analysis tool – it doesn't actually run the code through PHP and verify its assumptions are correct. The code analyser at this stage is a tool that will warn you about things that *might* be wrong – in most cases, at least in our tests, it does the job properly, but it's not yet any replacement for you and I at this point. You still need to check what it says is correct, then act upon it.

LXF: How many of the new features were inspired by feedback from your users? Do you get a lot of feature requests?

ZS: Absolutely – lots of it has come from feedback from our users. Some of the features are ideas that we have internally, but most of the key features are requests that come back from our users, which is great. The profiler came back as the number one requested feature from our users, so we went ahead and implemented it. The major features we implement always come from our users – previously the most requested feature was CVS support, and we implemented that, too.

LXF: Which features did you come up with internally?

ZS: The code analyser is something entirely new, as is the browser integration. The browser integration sounds like a small thing, but it has the potential to significantly change the way people work – we didn't get that as a request, but we wanted to come up with a way to help people switch from a primitive way of debugging (using *printf* or an *echo*) to using a real debugger.



“We wanted to reduce complexity of debugging to ‘trivial’ and I think we did a good job – debugging even the most complicated form is now just one click away.”

ZEEV SURASKI, CTO ZEND

The most difficult thing to change in this world is a habit – when people are used to debugging using echos, it's very, very difficult to change their habits, so we wanted to reduce the complexity of debugging to ‘trivial’, and I think we did a good job – debugging even the most complicated form is now just one click away. It's a small feature, but it has good potential to increase the level of acceptance of debugging in the PHP world.

LXF: And the browser bar should work on Linux?

ZS: It should work fine on Linux – we started on *IE*, but we're also working on a Mozilla/Netscape port. It may not have all the features of the *IE* version, but the main features – debugging and profiling from right from within the browser – that's going to be in.

LXF: Do you still think Java was a good choice for the product?

ZS: I can tell you both Andi [Andi Gutmans] and I were really not sure that Java was up to the task – we didn't think it could handle the job properly, and we weren't too convinced that it would be as easily portable as people said. We didn't really believe you could develop it on one platform and have it work out of the box on another platform. But the fact that we have to do very little work... to spend almost no time on porting, and get so many platforms out of the box – even platforms I don't think other people are looking at. People run it on Mac OS X – I even know people who run it on Solaris!

LXF: What's the release schedule?

ZS: As long as our Quality Assurance team do their job, the beta should be out by the time this gets to your readers, and that should help us find any big problems that our QA team missed. From there, we're hoping to have the final release by September 2003, although October is possible if we get pushed back – you never really know what's going to happen.

LXF: But isn't PHP 5 due out some time in February?

ZS: Actually, we still hope to release it by the end of the year – it's a bit of a self-fulfilling prophecy. You know, if I say it enough times maybe it will actually happen! We already released beta 1, and there weren't a huge number of bugs – there was a good amount of reports. That's good because it means people were trying it out, but it wasn't like “Damn, we're at least two years away.” The second beta should be out by the time readers have this. If the final version isn't out by the end of the year, it'll be very early next year – that's the latest we'd like to see it released by.

LXF: So how will you implement ZDE support?

ZS: Well, if there are any last-minute changes in the PHP 5 language syntax, then we will of course release a patch for *Zend Studio* to bring it up to date. However, if you look at previous PHP releases, by the time it reaches beta there weren't major changes to the language. We're working hard to help companies migrate to PHP 5 – as we're coming out with tools that support PHP 5 even before it's released, this should help people upgrade easily. ■

COVER FEATURE **DOMAIN NAMES**



Who's made off your . domain name?

Intellectual Property lawyer **DAVID HARRIS** explains the murky issues surrounding domain names and what some seemingly arbitrary ways in which the Web is governed mean for you and your business.

Domain names are yet another area of IP law that seems to be a common sense-free zone. In what was a novel area, tech-savvy entrepreneurs sought to make money out of lumbering corporates. While some of their demands were excessive and needed to be reined in, there has been little critical analysis of the real necessity for the resulting regulations and whether they strike the right balance. Does Megacorp really *have* to have Megacorp.com? With appropriate marketing and Google adverts they could function equally well with just Megacorpwidgets.com and Miss Meg A. Corp who got her domain first should be able to retain it, not just be bullied and intimidated into surrendering it by a multi-billion dollar company. Nonetheless, in an era where most big businesses seemingly continue to abide by the maxim “might is right”, trademark law is being abused and as with copyright what is at root a fair and proper system can be abused by one section of the community.

The Internet domain name system was originally run by the US government, various contractors and the National Science Foundation. In response to the exponential growth of the Internet, problems began to arise. For example, there was a trend for trademark holders to grab increasing numbers of .com domains, which created a pressure for new top level domains such as .biz, .shop, .name and the like. Additionally UK and European governments, ever more aware of the crucial value of the Internet to their economies, began to be more concerned about the US hegemony on domain name control. These political technical and infrastructure requirements resulted in the US Department of Commerce (DoC) devolving DNS management and associated decision-making ultimately to the Internet Corporation for Assigned Names and Numbers (ICANN).

US-centric

This has been a decision leading to controversy. The accusation made is that ICANN is being unlawfully used by the DoC either to delegate Internet policy or as an unlawful proxy for decision making that it should properly retain as a democratically accountable body. While affecting independence from DoC, ICANN does in reality launder hidden policy decisions of DoC while not in itself being open to any sort of genuine accountability or control; in fact even the illusory systems of elections at large are being removed and ‘temporary’ directors are seizing permanent control. This is likely to get worse if ICANN gets full control of the root rather than its discretionary loan from DoC. Control over domain names and routing facilitates the power to coerce name holders and applicants to conform to a potentially arbitrary set of controls. At the moment these are confined to mostly trademarks, an issue which we shall discuss shortly, but there is no reason why ICANN could not subject websites to other acceptability criteria. They could for example refuse or remove registration for sites that ‘promote terror’ or which have an unacceptable sexual content.

Control of the root currently lies, in practise, with the DoC. The US owns 7 of the 13 root servers and only 3 lie outside the US. There is some debate as to what the nature

DISCLAIMER

NOTHING IN THIS ARTICLE is intended as or should be construed as advice or acted on without seeking your own lawyers’ advice. The author cannot give personal legal advice.

of the legal control the DoC has over the root file in the A server; from which all the other root servers draw their data. It seems unlikely from its technical features that there is any form of copyright in it: it lacks originality being merely a collection of bare records, and there probably isn’t a compilation copyright. Notwithstanding that, the DoC does not seem willing to risk losing control of the root file, and it threatened John Postel with criminal prosecution when he attempted to move root away from the A server. DoC will also appreciate that few registries would risk the anarchy of a split root.

While such explicitly censorious policies are not yet anticipated, other equally destructive policies have been implemented. A system of compulsory arbitration in regard to trademark owners is imposed on registrants though the benefit or justice is unclear; it favours the non-incumbent whilst the registrant is subject to expense at the whim of an unlimited number of third parties. The legal basis often seems to be barely connected to the trademark law it pretends to support.

What’s in a name?

When the Internet began its growth spurt in the mid 1990s, the system began to creak as the *ad-hoc* administration manning the naming system struggled to keep up with the demand. Part of the reason for this demand was that the legal analysis of domain names was unclear. Lawyers like me advised businesses to grab as many of the domain names that were related to their business as possible – for safety’s sake. To further compound the rush entrepreneurs, thesaurus in hand, began the great land grab for good generic domains motivated by nonsense like the sale of business.com for \$7.5 million in 1999.

Trademarks, passing off and unfair competition

Disputes over domain names revolve around trademark law, since it is this facet of IP law that governs the commercial use of names and signs. Trademarks are a pivotal part of intellectual property law for businesses. They are the thing

“ICANN does in reality launder hidden policy decisions of DoC while not in itself being open to accountability or control.”

that gives each company its uniqueness and commercial value. Understandably therefore, most corporations guard trademarks carefully from abuse or degradation and the risk of the latter from their use as domain names did not escape the attention of lawyers for long.

Part of the clash between the interests of trademark holders and domain holders lies in the fact that domain names need to be unique: there can be only one www.linuxformat.co.uk. A trademark however is not subject to such a technical constraint, and by commercial tradition they have been applied to distinct and different classes of goods. These classes of goods are defined by the *Madrid Protocol* and trademark



COVER FEATURE **DOMAIN NAMES**

owners will usually, except for famous marks, be able to register a trademark only within those classes they use. 'Acme' for example is a mark that might be used by several different traders for goods in unrelated areas of commerce, *eg* Lawnmowers and photocopiers. Under the domain name system Acme Lawnmowers Inc. will be compelled to argue with Acme Photocopiers Inc. over who has the right to acme.com, and the policy of most domain name registrars is that, between equally entitled applicants, there is a 'first come, first served' system of registration. This view has been supported in the UK by the *Pitman* case between Pitman Publishing and Pitman Training. These two unconnected companies have trademarks in their own class and they have traded in parallel for decades without clashing. Pitman Publishing got the domain name but didn't use it. When it got round to trying to use it it found Pitman Training had, because of a (supposedly impossible) error by Nominet, who re-registered it to Pitman Training a month later. Pitman Publishing got the domain transferred back, claiming it registered it first and was entitled to it on the first come, first served basis of registrar Nominet. Litigation ensued and the judge held there was no evidence of passing off; further, both companies had traded under the same name for years, and passing off was not a viable cause of action, but just a consequence of the valid use of parallel trademarks. Pitman Publishing registered the domain name first, so were entitled to retain it.


Trademarks defined

A trademark is, broadly, a sign that can be used to distinguish the goods or services of one trader from those of another. Alternative definitions exist whereby the trademark only serves to designate the quality of the goods or services on which the trademark is applied. The basis being that consumers often don't know or care who makes something, but are only interested in whether its any good or not. Linus Torvalds, owner of the 'Linux' trademark, might well be entitled to prevent Microsoft from marketing 'MicrosoftLinux' because of the confusion over the origin and quality of the goods, *ie* the operating system. Pertinent to this discussion is what a 'sign' is; it is not fully statutorily defined but statute and case law gives examples. Clearly

"Generic trademarks like 'banana' cannot be registered, while generic domain names can be, and are often highly desirable."

written words and/or pictures function as a sign but also colours, smells, shapes, slogans and sounds can be trademarks if they can also be graphically represented, such as by musical notations. Likewise shapes such as the Coke bottle can be trademarks.

Companies occasionally want to register generic or highly descriptive marks but that is usually impossible, or at least very difficult without it being well known as connected with the applicant. If the function of a trademark is to distinguish the origin of goods, a mark like 'BANANA' should remain



available to any trader who wants it, and this helps to prevent any single trader acquire a monopoly. In reality, marks lie on a scale that ranges from the fanciful, to the suggestive to the generic. Purely generic marks like banana cannot be registered, whilst generic domain names can be registered and are often highly desirable. This may in part explain the attempts at reverse hijacking of domains: if a valuable but unregistrable mark cannot be trademarked, getting it as a domain name is a good option. One must be cautious however about what a generic mark is; some marks which are said to be generic become less so if as a result of actual use, people come to recognise them as associated with a particular company.

Geographical marks will generally not be registrable unless there is strong evidence that the public associate the geographical name with a particular trader. Likewise, some well-recognised symbols cannot be trademarked: national flags and the emblems of some international organisations *eg* the UN or the Olympics. The names of royalty and US presidents cannot be trademarked either.

Surnames cannot generally be trademarked where they are primarily mere surnames, since many traders would want to use their own name, and additionally the public won't associate it with any one trader. Where they do, trademarking may be possible; although the Estate of Elvis Presley was unable to persuade the UK's High Court they should have it as a trademark since the court viewed the public as uninterested in where the goods came from or if they were "genuine".

In passing, we should also mention the concept of famous marks and trademark dilution. Since trademarks differentiate product classes, two identical signs in different classes are permissible since no one will be confused. IBM computers are unlikely to be confused by anyone with IBM cooked foods. However, owners of famous marks have bridled at this and they say that no one should take a free ride on their mark in a way that may lead to its value being diluted by such otherwise lawful use. This is the position in the case of the Mobilix website, which discusses mobile Unix, which has had mixed fortunes in court against Les Éditions Albert René, the proprietor of the Obelix trademark famous from the world-renowned Asterix cartoon series. The allegation has been that Mobilix have taken unfair advantage of the similarity between Mobilix and Obelix. Of course the naming of Mobilix derives from no more than the collation of the words mobile and Unix. The German Court of Appeal overturned a lower court decision and said that the marks were indeed confusingly similar. This seems, frankly, idiotic on the limited facts so far published; in trademark dilution cases there needs to be more than just similarity but also unfairness in the usage and confusion between the two uses: in the UK in the Oasis Stores trademark application the court said that it was not the intent of the law relating to famous marks to prevent registration of all similar marks but only those where some confusion would arise. The position under German law must be similar, since UK and German law arise from the same EU directive. How one can confuse a cartoon character with a computer operating system is not entirely clear, but the German Court of Appeal seems to believe it's

possible. The hope remains that a final appeal will render common sense.

There is thus no doubt that a domain name is a 'sign' under trademark law, since it is composed of written characters and can designate the origin of goods or services. Both the US and UK and EU patent offices will register domain names as trademarks.

Passing off/unfair competition

I use the two terms interchangeably in this article – passing off is broadly equivalent to 'unfair competition' in America – both are remedies related to trademarks that often arise in domain name cases. Passing off usually arises when one trader sells goods or services in such a way that the people buying them believe they are coming from someone else. If some dodgy 'Del Boy' character sells watches labelled 'Roleks' but which look like Rolex watches, then an action for passing off may lie at the suit of the Rolex trademark holder if buyers are confused.

The essence therefore is a calculated misrepresentation of the origin of goods, which causes confusion and damage to another trader. Although I talk of a 'calculated misrepresentation' it is not necessary for this confusion to be deliberate – an innocent trader can be sued for passing off also, and the phraseology is confusing to the non-lawyer. What amounts to misrepresentation is regarded flexibly by courts: the overall impression of the style and get-up of the product is assessed. Thus a mis-spelt domain name could be taken to be a passing off. In one case, Microsoft were able to get the domain name MicrosOft in an allegation of unfair competition against Zero Micro. The basis being that the impression was one of a connection between Microsoft and Zero Micro, despite the fact that on the face of it, MicrosOft is a rather good name for Zero Micro (though microsOft.com currently appears to be owned by someone other than Microsoft). It need not be just names that are complained of, it can be the general impression. Easygroup has become infamous for its belief that it owns the commercial use of the word 'easy' and has used passing off law to claim that people acquiring domain names with the word in it are passing themselves off as being connected with the Easygroup. In at least one case the fact that a website used similar colours styles and fonts to Easygroup enabled them to acquire a domain but in many others there is palpably no connection or confusion, and the courts have sent Easygroup packing (see *Resources* box overpage).

Trademark and passing off defenses

Defenses exist to trademark infringements that are germane to any debate on domain names. A common theme in the defenses is that an infringement has to compromise some exclusive commercial interest since it is this that which goes to the root of the trademark function. Defenses to trademark infringement will typically incorporate the following elements;

■ **Use of ones own name or address or registered trademark.** In the UK the *Trademarks Act 1994* (TMA) provides that honest use of ones own name or address is a defence and in the US this would be fair dealing under the § 33(b)(4) of the *Lanham Act*.

REVERSE DOMAIN NAME HIJACKING

Suit you, sir? Or would sir prefer a LAWSuit?

ARMANI CLOTHES WANTED A WEBSITE

Armani.com but unfortunately for them, Mr Mani got there first. Mr Anand Ramnath Mani that is. He was quick, they were slow. What does an Italian couturier with a famous trademark do in this tricky situation? A corporate mugging is the normal form, but Mani had irrefutable evidence to show *bona fide* use of his name in a trading style since the 1980s. Armani suggested that the confusion caused by his use of his initials meant he had no legitimate interest in the domain. The panel said that he was entitled to use his initials and he did not have to use a non-confusing variant such as anandmani.com. The panel suggested that Armani acted in an abuse of process by trying to reverse-hijack the domain merely because they had a famous mark. The \$2000 that Mr Mani asked for was a very fair sum at that time and no more than his costs.

Famous people

We have mentioned in the main body of this article the rockstar Sting (formerly of the highly successful 1980s band The Police) who was unable to get sting.com. This was a generic phrase and the registrant made honest use of it prior to Sting's complaint. By contrast Madonna, though held as a trademark by people other than the singer, and having many distinct meanings, was registered unfairly on the facts. The registrant had apparently intended to use it for a sex site, so the decision was probably correct but had some very worrying elements such as the panel declaring it didn't feel bound to respect a foreign trademark registration when deciding in favour of a US complainant if the foreign trademark was not obtained in 'good faith'. Nor, surprisingly, was an obvious disclaimer regarded as enough to dispel confusion.

■ **Parody.** In America, it was said that the Muppets pig character Sp'am did not infringe Hormel's trademark in the meat product 'Spam'. In the UK it was acceptable to do an S&M parody of the Miss World contest called 'Miss Alternative World'. In neither case was there said to be an infringement due to any confusion over origin; the public would know what was false and what real.

There are limits however and in America these seem to be morally based: "Enjoy Cocaine" violated Coca-Cola's slogan "Enjoy Coca-Cola". The T-shirt "AdiHash" displaying a cannabis leaf that parodies the Adidas logo would seem to be similarly vulnerable in the US but perhaps not the UK.

■ **Mere designation of origin.** In the US the doctrine known as 'nominative use' permits the mention of a trademark if the purpose is to show where goods come from. The UK has a similar rule. There is no rule that says only the trademark owner can use their own trademark. If I wish to mention the Apple iMac OS X machine in this article I do not infringe the trademarks 'Apple' 'iMac' or 'OS X' since all I am doing is designating their goods for some fair non-competing purpose such as commentary.

■ **Descriptive use.** Where honest use of a trademark is made just to tell users the kind, quantity, quality, purpose or geographic origin of goods there is a defence both under the *Lanham Act* and the TMA. This provision stops people who have obtained a borderline trademark from abstracting common material from the public domain. We have already mentioned the controversial Easygroup who appear to believe they own the word 'easy'. Despite much litigation against honest traders, they have been consistently defeated in court because people who use the word easy have been doing no more than making an indication of the quality or kind of the intended

COVER FEATURE DOMAIN NAMES

◀ goods or services (see the link in the *Resources* box below for more info).

■ **Indication of the intended use.** This defence is primarily aimed at the situation where spare parts or product related services are needed. The intent is to promote competition by ensuring that it is not just the trademark owner who can supply spares and servicing. Were it otherwise, manufacturers could lock in purchasers to their own service agents: no-one else would be permitted effective advertising. On this basis a domain name such as ford-cheapcarservicing.co.uk might well be permissible and in fact this was the result of a case called *BMW v Deenik*.

Domain names and trademarks

Some difficulties soon became apparent when lawyers, steeped in the above law, began to notice the similarity of a trademark to a domain name: there was a debate about what a domain name was in law and whether it was amenable to trademark law at all. John Postel wrote initially in RFC 1591:

"The registration of a domain name does not have any Trademark status. It is up to the requestor to be sure he is not violating anyone else's Trademark".

Thus the naive intent of the fathers of the domain name system is made clear; domain names and trademarks are separate and the managers of the system did not want to be involved in any disputes.

Some commentators wondered if perhaps Postel was right; wasn't a domain like somewhat like a street address? Could it not direct you to a virtual shop? Just as a street address or country cannot usually be trademarked, neither normally should a domain name. This theory was soon rejected in the courts of most countries: in the American case *MTV Networks v Curry*, the judge said that a domain name was more like a phone number mnemonic such as 1-800-COLLECT where the letters correspond to numbers on the pad. The mnemonic numbers are, unlike addresses, chosen by the user. These cases have been largely confined to the US where freephone numbers are more widely used than elsewhere. In the main English case on domain names 'One in a Million', the Court of Appeal viewed the registration of bt.org, virgin.org and other domain names as the registration of a sign the use of which was capable of infringing a registered mark. On both sides of the Atlantic



the result was clearly that courts would not regard domain names as unique but as something capable of treatment under trademark law or related laws such as passing off.

The initial attitude of NSI was to say that it was not itself responsible; it didn't do trademark searches when creating domain names and merely doled them out on a 'first come, first served' basis. It was not its problem and it merely stood between the parties. It changed policy after KnowledgeNet sued it, and it began to permit trademark owners to retrieve domain names. To make that job easier, it also instituted the privacy-busting requirement of people publishing their names and addresses, since this made the job of trademark owners tracking down registrants easier, a controversial policy which remains to this day.

Domain name legislation

In the US, domain names that incorporate trademarks are subject to the Anticybersquatting Consumer Protection Act. This Act was introduced when domain 'hijacking' was beginning to hit the headlines. The Act largely parrots ICANN's *Uniform Domain Name Dispute Resolution Policy* (UDRP) and is therefore just as flawed. It specifies factors to be considered in "bad faith" registrations of marks which are nearly identical to the UDRP; two examples are:

- the registrant's provision of material and misleading false contact information when applying for the registration of the domain name;
- the registrant's intent to divert consumers from the trademark owner's online location to a site accessible under the domain name that could harm the goodwill represented by the trademark; and providing invalid contact information should not be a factor in a domain dispute.

The intent is clearly to target those reprobates intending to evade litigation, however courts have effective and long-established methods of suing people who refuse to acknowledge service of process or are not contactable. More important is the need to protect privacy: if I want to set up a web site like abortionisawomansright.org, I might, justifiably, worry about a knock on the door in the middle of the night followed by a bullet. These worries can be circumvented, but why should one have to jump through a lot of seemingly unnecessary hoops in order to protect one's right to free speech?

The Act also introduces some rather spurious protection for the names of individuals. This portion of the Act was drafted after the lobbying of famous people some of whom have not been able to get their stage names as domains because fans or speculators got their first. For example Sting was unable to get the domain sting.com given that it was both a nickname, generic and had been legitimately used by its registrant.

A concern over the Act is that it fails to consider non-US interests. For example it potentially applies to domains that are registered by national registries. This could lead to a US corporation claiming the right under the Act to acquire a .co.uk domain. Nor is there any consideration of the issue where two people are equally entitled to the use of a trademark in different countries; are the US courts going to give priority to US trademarks? Other portions of the act are potentially fair and should function well

RESOURCES

<http://www.icannwatch.org/>
<http://www.mama-tech.com/antipiracy.html>
<http://ecommerce.wipo.int/domains/>
<http://www.easyprotest2.com/>
<http://personal.law.miami.edu/~froomkin/articles/icann-main.htm>
<http://www.cavebear.com/cbblog/> – Karl Auerbach's blog

UDRP

<http://www.icann.org/dndr/udrp/policy.htm>
<http://www.nominet.org.uk/DisputeResolution/DrsPolicy/DrsPolicy.html>
<http://www.udrpinfo.com/> – Prof Michael Geist's web page

PROTECTING YOUR FOSS DOMAIN

Geeks v domain grabbers

ITS WORTHWHILE TO NOTE THAT THE UDRP policy does not require you to have a trademark in order to protect your domain name; it is merely necessary to have a legitimate interest under 4(a)(ii). Whilst it is best to have a trademark, only the richest of geeks can afford these at several hundred dollars each. A moderate reputation for the project will still give it protectable rights at common law under the rules of passing off/unfair competition. These rights can be used to defeat a domain grabber. In theory the unregistered trademark rights should even prevent a registered trademark holder who

also wants the reverse hijack domain.

It should be noted that many FOSS projects are not commercial and the code is free beer. One potential difficulty is that since there is no trade, there might not be any right to a trademark. In reality this may not be an issue for a lot, but not all, projects. In Europe and the US, trade has been fairly liberally defined by the courts; charities, political parties and even religions (and what some see as quasi-religions like Scientology) have been deemed to be in trade. Further, in the EU there is no explicit requirement in the legislation requiring trade, it speaks merely of undertakings which

is a very loose definition. On this basis those projects which are sponsored, either commercially or by a cash equivalent such as free server space, or which solicit funding will be able to obtain a trademark. Some projects might however have difficulty; for example the lone hacker or those where this is no hint of actual or potential commerce.

In all cases, having a reputation is important and for this it would be useful to have a Sourceforge listing which is active or a Freshmeat entry. This will help to establish a reputation in the market and that is a factor that weighs heavily with judges and panellists.

provided they are not applied over broadly by the courts in favour of corporations.

ICANN

NSI gave way to ICANN with the hope that it would be a more effective and businesslike domain name administration. In many respects it has succeeded but has in the process become entirely controversial (for a sample of which, www.icannwatch.org is an invaluable guide). It has a hidden, over-close relationship with the US department of commerce and its partiality to trademark owners is a matter we discuss below.

UDRP – dispute resolution

ICANN has taken the approach of mandatory arbitration as a means of resolving disputes between trademark owners and domain name holders. It has the advantage of being slightly quicker and cheaper than litigating in court, and because of it ICANN has less risk being a party to proceedings. There are many problems with the UDRP which were hoped to be initial teething problems but which appear to have become somewhat embedded. Historically domain name disputes featured so called 'domain hijacking' by speculators looking to acquire valuable names that they could sell to the trademark holders. This seems to have impressed ICANN's policy and it has not sufficiently taken into account issues of fair use such as free speech and criticism. Nor does it combat the more cynical behaviour of some corporations who see a valuable domain which is taken and who then try to acquire it by litigation in process known as a 'reverse' domain name hijack. This was not much of an issue when the UDRP was designed but ICANN have displayed singular ineptness in designing procedures to combat it, leaving it to individual arbitrators to handle.

Nor is the procedure for arbitrating ideal. ICANN has two bodies it uses for arbitration: WIPO and the National Arbitration Forum (NAF). These two is responsible for selecting arbitrators and did the initial drafts of the UDRP however their initial draft of it *excluded* fair use! This underlines one of the more persistent criticisms of WIPO – that it is a corporate club with no public interest

representation or indeed interest. Much of the more controversial amendments in global intellectual property law have had their genesis at WIPO, so it is perhaps not surprising that they did not consider fair use to be important.

For most non-UK disputes, the ICANN dispute rules apply. In the UK, a very similar policy designed by UK registrar Nominet exists, but with some important distinctions. In summary the UDRP rules provide that domain registrations can be amended on the basis of a court order or legal requirement or an arbitration panel decision. The arbitration is compulsory and is based on three considerations:

- 1 an assertion is alleged of trademark infringement; and
 - 2 that there is no legitimate interest in the domain; and
 - 3 the registration has been made in bad faith
- Bad faith features highly in complaints and ICANN specify factors to be assessed in determining it:

“ICANN’s policy does not sufficiently take into account issues of fair use, nor combat cynical behaviour of some corporations.”

- i) it was obtained for reselling to the trademark owner or a competitor for a sum that is more than out of pocket expenses; or
- ii) to prevent the domain owner from using his trademark in a domain name but only if there is a pattern of such behaviour; or
- iii) for disrupting the business of a competitor; or
- iv) for the purposes of diverting surfers to a web site by causing confusing as to source affiliation or sponsorship of a site or its contents.

Arbitration – ICANN can I?

The UDRP has caused some controversy over the mechanisms by which panellists are selected and by whom, and how they reach decisions. There are accusations that forum shopping is occurring and panellist selection results in undue advantage to trademark owners.



COVER FEATURE DOMAIN NAMES



Arbitrators, or panellists as they are known, are not employees of ICANN, WIPO or NAF. They tend to be lawyers academics or similarly experienced individuals. Their fee for acting as arbitrator is between \$1000 and \$1750 for what is a not terribly difficult or lengthy piece of work. WIPO or NAF take a large slice of the fee, so they have a substantial commercial interest in being popular. It's nice work if you can get it. The accusation is that whether you get it depends on your reputation among those most likely to pay the fee or give you the case to decide: that is trademark holders and the body providing arbitration, WIPO or NAF; there is therefore the likelihood of partiality. The overwhelming number of complainants are

“UDRP arbitration is an administrative and not a legal process: precedents have no place, nor do the rules consider jurisdiction.”

trademark holders; a non-trademark holder cannot make a complaint against a trademark holder or make a request for a determination; for example, if he wishes to determine the position before investing lots of money in a site.

I take the view that there is not usually explicit bias by arbitrators, but that the terms of the UDRP together with

the professional experience of arbitrators tends to give an unconscious sympathy for trademark holders. Some arbitrators have no experience in trademarks law and this leads to some truly idiotic decisions, discussed later, that wouldn't be made by a first-year law student. In any event, some of the trademark principles applied don't conform to any known principle under common law and other conflict with it directly. For example, 'bad faith' has a meaning that would be intuitively understood by most lawyers as implying the application of certain equitable rules based on elements such as honesty, best efforts, faithfulness and impartiality. To some arbitrators, it appears to have become a catch-all class of 'things we don't really like'. This tendency to trademark holder sympathy is manipulated by WIPO and NAF, and in a secret and non-transparent process when they can they load the panels with complainant-friendly panellists. It's remarkable that of the 130 panellists, 50% of cases are dealt with by six people (according to a study by Prof. Michael Geist – see *Resources* box) and 90% of cases with a single panellist are resolved in favour of the complainant. Only in 3 panel decisions were the odds of complainants winning reduced to 60% and in these cases the influence of NAF and WIPO are reduced by the ability of both sides to influence the choice of panellist.

ICANN has defended the process by saying that the difference is caused by the fact that respondents fail to react to the complaint and that therefore they tend to lose their case. However follow-up studies (see *Resources*) indicate that this doesn't make any difference and the percentage figures remain just as skewed when this is taken into account.

The UDRP arbitration process is an administrative and not a legal process, in which case precedents have no place. Arbitrators therefore often make different decisions on virtually identical facts since they are not bound to consider previous cases. Nor do the rules provide any basis for considering jurisdiction: a French complainant may apply for a decision against a Brazilian domain owner and have a hearing made on the basis of a bastardised version of US trademark law. The result appears to be that too often people who have little understanding of trademark law make decisions that are inconsistent with each other. The points made above are best illustrated by considering some of the higher profile arbitrations. This is only a small flavour of cases and I'd refer readers to the WIPO site (see *Resources* box) for further examples.

“Suck” domains

“Suck” domains seem to generate a lot of heat in the arbitration process. Most people would say that only a moron in a hurry would regard a suck domain as capable of being confused with the offerings of the complainant. Some panellists take the same view and this is an area where one sees the rarity of numerous dissenting opinions. There is even a WIPO sucks domain (which is not being put to the best possible use at the time of writing).

Walmart v Harvey

Mr Harvey is a Canadian who had a store called Walmarket and he got the domain name walmart.com but lost it to the US company Walmart after offering to sell it to them for \$5 million. This led to some ill feeling and he began registering

HOW TO LOSE A DOMAIN

Some features become clear indicators of impending loss of a domain:

Offer to sell a domain to a trademark holder or his competitor.

A common trick is for the holder to initiate discussions about the sale of the domain. Once this happens they will cite it as bad faith even though they began negotiations!

Offers to sell for legitimate transfer costs are ok, that is several hundred dollars or maybe a thousand or so may be ok depending on the circumstances.

Register multiple domains incorporating other peoples trademarks.

Just registering multiple domains seems not to be problematic *per se*.

Make a use of a trademark name in a confusing fashion.

For example offering goods or services that compete with the trademark holders or which cause confusion. For example auto-redirection to an unconnected competing site.

Register misspellings or hyphenated versions of trademarked names; eg MircOsoft.com Or 1bm.com (though neither of these are held by their

trademark owners I contend it would be trivial to get them transferred).

Fail to file a response to a complaint.

The majority, though by no means all, of uncontested hearings result in a transfer decision being made.

Fail to select the right panel

Statistically complainants usually win if they select a single panellist whilst respondents have a better chance if they get a three-member panel.

How to defend a domain

Given the dismal record of WIPO, an accused domain holder will usually be better off initiating litigation in a national court despite the extra expense. Many of the more dubious assertions made by panellists have been contradicted by courts. However, given the permissive attitude of the US courts and trademark registries to trademarks, one is better off not litigating in the US but in Europe or Australia. Of course since domains are so cheap, serious consideration should be given to just walking away from it.

large numbers of "Walmart suck" domains. Walmart argued that the plan to disparage Walmart would damage them and denied he had absolute freedom of speech to criticise them. Walmart were successful and the suck domains were transferred to Walmart. This was however despite a court case concerning luentsucks.com, delivering the common sense conclusion that consumers would *not* be confused because of the addition of the suck suffix. The panellist seemed to be swayed by the unattractive visage of Mr Harvey into twisting legal logic to deliver what it perceived as justice rather than an intellectually rigorous decision.

However, there are rare cases of good sense; in the wallmartcanadasucks.com case, Walmart lost on very similar facts with the panellist deciding sucks was a legitimate form of criticism and parody that did not amount to a confusing similarity of marks. The panellist Henry H. Perritt made the eminently sensible comment: "...distasteful conduct should not stampede UDRP decision makers into an unwarranted expansion of the domain name dispute process..." He said that the respondent was unfair, childish, retaliatory and indulging in unwarranted disparagement but this was not behaviour that the UDRP could or should redress.

American court cases such as *People for the Ethical Treatment of Animals v. Doughney* have stressed the need for the domain name alone to make it clear that it is a parody and not only once the visitor has arrived at and read the site. Peta is a vegetarian anti-meat-eating site so [petasucks](http://petasucks.com) justeatmeat.org would be ok but peta.org isn't.

Generic/descriptive domains

Perhaps a bit of surprise to many was the decision on the unix.org domain. This illustrates the danger that often what is thought of as a generic domain name is not. Unix is the trademark of the X/Open company while the registrant believed (as many do) that Unix is generic. In part it was due to their low profile, and the respondent did argue that the mark had become generic through a lack of policing and common use in the trade. However the arbitrator made the point that a trademark registration raised a *prima facie* case of validity which the panel could not decide on since it was not a court. The balance of the decision had some rather peculiar assertions: mere non-use of the domain was an act of bad faith; as was adding links to other commercial Unix sites. The rest of the decision is hopeless: arguments which are supposed to be independent are treated as dependent on each other and conflated; there has use of a trademark and the respondent must not have a legitimate interest in the domain, but here it was said that they had used the trademark and therefore had no legitimate use for the Unix domain. This is not the proper application of the tests and is not untypical of panellist decisions. The domain was ordered to be transferred to the X/Open company.

There seems to be no end of hopeless generic domain cases and WIPO seems, uncharacteristically, to deny most of them but one has to wonder what these complainants legal advisers are telling them. Attempts have been made, for example, to reverse hijack internetcontent.com, snooze.com and even rollerblading.com, all of which were rightly rejected. In no small part the fault here lies with the cretinous policies of the US Trademark Office



which often allow trademarks that are very close to generic, based on minimal commercial usage and differentiation. This has led many a registrant to believe it has rights in generic names and to the cases before panellists.

Domain seizure in criminal cases

Civil and administrative processes are not the only means by which a domain can be lost. American police have deployed a new technique for restraining and investigating crime: seizure of domains. Federal courts in the US are authorising the police to seize domains and point http requests at their web servers in cases where crime is alleged.

Many countries have laws permitting the seizure of property in criminal cases or closely related matters. For example in Canada the *Controlled Drugs and Substances Act* or the *Proceeds of Crime Act* in the UK, and civil asset forfeiture has a long and controversial history in the US. It was only a matter of time before someone joined up the dots and figured out that domains were a species of property and hence subject to seizure under numerous legislative provisions like these. isonews.com was a warez site that had previously been threatened by Sony and Sega. Recently however warez d00ds looking to go to the forums to trade discovered that it now points to an American Department of Justice webserver which provides a polemic on the evils of warez and copyright infringement. The site had been surrendered pursuant to a plea bargain related to charges laid as a result of selling Xbox modchips. The same procedure has been followed without convictions or even criminal charges being laid in cases related to drugs paraphernalia: [PipesForYou.com](http://pipesforyou.com) and several similar sites have also been grabbed. There are of course privacy concerns, since visitors to these sites are having their IP addresses recorded not by private operators, but by police agencies; this raises civil liberties concerns among privacy advocates at actions that may amount to either unlawful surveillance or potentially entrapment.

In addition, there are international dimensions, since the global reach of web sites may result in police seizing websites that relate to activity that is lawful elsewhere. For example, a Dutch file-sharing or cannabis site in the Netherlands might be targeted by US police on the basis that it can be viewed within the US. Indeed, Italian police used just such a justification for altering without the consent of the owner servers based in the US that were distributing 'blasphemous' content into Italy; this could potentially lead to the FBI arresting Italian detectives for cracking.

Conclusion

The Internet domain name system is another area in which intellectual property laws are seen by many as being deployed to the benefit of corporations at the expense of freedom of expression and often common sense. Structures have been set in place that, as with copyright and patents, take too little interest in the rights of the public and pander to one sector of the community at the expense of others. Arguably the damaging effects of this are less serious than those arising due to the DMCA or EUCD. Nonetheless, it is indicative of a growing trend in legislatures and and governance bodies that needs to be combated. ■

FLEX AND BISON



flex and bison. Compiler writing

Despite me promising to move onto how intermediate code is a big step forward, there's been a minor change of plan – we've been inundated with letters from readers asking various questions about compiler creation, from fairly simple things we covered in the first issue to more advanced topics that are yet to come in this series. In order to make sure everyone's getting the most out of these tutorials, I want to answer these questions before continuing. So, without further ado...

Q Is SKYLang platform-specific? I run Mac OS and/or Windows as well as Linux, and want to try my language on other platforms.

A To create SKYLang, the reference language behind this tutorial, we've used *Flex*, *Bison*, and GCC – all tools available on many platforms. So long as you've used vanilla C++, that is, you haven't called any operating system-specific functions, you should be able to take your *Flex* and *Bison* code and recompile it smoothly for any other platform. On Windows, for example, you can download *Flex*, *Bison*, and GCC as part of *Cygwin*, and SKYLang will compile smoothly there. If you want a more native port for Windows, you'll need to find a C++ compiler plus a non-*Cygwin* version of *Flex* and *Bison*.

Q Why is intermediate code necessary?

A Bison processes its source scripts by using *Flex* to read match tokens as it proceeds through a file, which means that *Flex* moves linearly through a SKYLang source script picking up tokens as it goes. Now, consider a conditional statement – for example,

```
if ($foo = $bar) { do_some_code(); }
```

If that evaluates to be false, we need to skip the conditional code block and continue executing immediately

PART FOUR
PAUL HUDSON
takes a brief hiatus from the intense world of compiler writing to answer a few of the deluge of questions we have received...

afterwards – in effect, wind the *Flex* source script forward till it gets to where we need to be. Also in statements such as `for $i = 1 to 10 increment 2 { do_some_code(); }`, we'd need to keep moving the *Flex* file pointer around each time we went through this loop.

Not only would that be incredibly slow to do, but it's also actually impossible – *Flex* does its reading linearly, with no way for us to override that. So, in order to properly handle these sorts of situations, we allow *Flex* to read the entire file through, then convert it to intermediate code – once it's in intermediate code, we can jump around as needed.

Q Why didn't you create a C++ interpreter for the reference?

A While creating a C++ interpreter would allow me to teach many more concepts than teaching SKYLang, it would need to have come in a 72-part series in order to get a working program! Also, while it would allow me to discuss how to handle pointers and references, it would limit the chance to discuss language design – a topic that's generally much more of interest to compiler newcomers.

Q How hard is it to make a new front-end?

A As of right now, our back-end is pretty flimsy – we iterate through op-codes and perform various operations. Our front-end is equally straightforward – much of the work is done by our lexer, with little bits here and there being done by our *Bison* parser. If we were to create a new language and wanted it to compile down to the same intermediate code as SKYLang with the goal of having both languages work with the same back-end, the only major changes would be the need for a new lexer plus new parser rules – we could use the same code to handle binary operations and such, as well as the exact same back-end. Perhaps this

is something we'll be looking at later on – do let me know your thoughts.

Q How can I make SKYLang execute its scripts faster?

A We'll be looking at optimising intermediate code in upcoming issues – it's quite a complicated topic, but hopefully I'll be able to make some sense! There are a number of basic optimisations you can try out before I cover them in-depth – strength reduction, for instance is where slow operations are replaced with faster ones, for example `"i = i * 2"` is equivalent to `"i = i + i"`, with the latter being a faster operation to execute.

An easy way to make things faster is to use the `-F` parameter to *Flex*, which makes it generate a fast lexer. You might wonder "why isn't this always set?", and the reason is because some other options don't work with it – *Yacc* compatibility, for instance.

Q Will optimised SKYLang be very fast?

A Probably not, no – at the end of the day, this is only a short tutorial series examining the basics of compiler writing. Whole books of theory have been written on this topic, and compilers such as GCC have been under development for years to eke out maximum performance. One of the goals I hope to accomplish with the tutorial is to make the language output C++ code that can then be compiled with GCC – that's likely to be the best way for SKYLang to execute scripts. Using this method, the SKYLang interpreter can then compare the C++ binary script file with its source SKYLang script file and execute the cached binary if it's newer, and it also means we're able to take advantage of all the hard work that's been done by the GCC team.

Q Why is SKYLang typeless?

A Ah, good question! The decision was ultimately made by the fact that PHP is typeless, and PHP is written using *Flex* and *Bison* – the "ultimate goal" of the tutorial is to give readers enough information about compilers that they should be able to read the PHP source code (or indeed any other well-written *Flex/Bison*-based language) and understand how it works, why they made certain choices, etc. The PHP implementation is remarkably clean and easy to understand, and so makes a very good step forward once this tutorial finishes.

Q Will SKYLang ever be typed?

A I've considered making SKYLang into a typed language near the end of the tutorial – not for any functional purpose in the code as such, but more because it allows us to teach readers about a whole new area of compiler writing, which is checking type safety for assignments and function parameters.

Q Will SKYLang ever be object-oriented?

A No.

Q Can I send you my patches to your code?

A Sure, go ahead! We can't guarantee anything will be printed in the magazine, but if you've got a cool hack

you've added to the SKYLang codebase, send it in with a brief description and we'll take it from there. Any particularly interesting additions may well be "made official" and merged in with the main SKYLang code. Why not share your thoughts and developments of SKYLang with other readers on the *LXF* forum too?

Q Will SKYLang be able to support existing libraries?

A Once we're through with intermediate code, the next step is to take a look at incorporating external libraries – there's such a wide variety out there that it will probably be a case of me picking one randomly out of a hat and implementing it in SKYLang. The goal is, of course, to be able to use the library functions direct from within SKYLang, which would add much more power to the language with very little work. While it would be ideal to cover how to implement an extension architecture, it's also out of the remit of this tutorial and is likely only to be covered from a theoretical point of view.

Q Why did you choose C++ to write SKYLang?

A I've been using C++ for quite some time, so it comes as second nature to me, however I was considering writing the whole thing in C at first because it has the extra edge in cross-platform compatibility. The main reason why I chose C++ is because it allows me to teach the concepts of compilation (how to store variables, how to generate intermediate code, etc) without having to spend time teaching data structures. Instead we've used STL structures like `std::map` and `std::vector` that do all the hard work for us, and are probably just as fast – as a result, I've had more time and space to cover the important topics.

Q What's the best book available on creating compilers?

A Most people will tell you that it's the infamous "Dragon Book", *Compilers: Principles, Techniques, and Tools* by Alfred Aho et al (Addison-Wesley ISBN: 0-2011-0088-6) – it's quite old, but covers compilation using hand-written lexers and parsers and also *lex/yacc* (the precursors to *Flex* and *Bison*), as well as intermediate code generation, code optimisation, and more.

However, I disagree – while it's a great book by all means, I would say that it has recently been knocked off the top spot by a newer book, *Advanced Compiler Design and Implementation* by Steve Muchnick (Morgan Kaufmann ISBN 1-5586-0320-4). Muchnick covers the topic in a brand new light – the latest techniques are discussed in depth, alongside more simple projects accessible to even less-experienced readers. One high point of the book is that over half of it covers optimisation – it's very comprehensive for readers wanting to maximise the performance of their compiler/interpreter.

Q Can SKYLang be converted to Java bytecode or .NET?

A The advantage to the front-end/back-end approach made possible by using intermediate code is that we can send anything we want as the output language, whether that be C++, x86 assembly, or .NET intermediate language. Actually, I rather like the idea of SKY.NET – I'll pass it onto my friend at Cyberdyne, Miles Dyson... ■

WHICH HOST?



Which web host is right for you?

Having a web presence is almost mandatory in today's business environment, and as a result there are a large number of companies waiting in the wings to provide you with that presence. With each customer having their own particular requirements for their web service, each web provider offers different solutions – some rely on low prices to woo customers, others rely on the quality of their support.

We reviewed six of the top web hosting companies in the UK, testing how fast their servers were, what they provided as standard, as well as how good their technical support is. Each of the web hosts were given just one requirement: their service must cost under £500 a month, which is about right for an enterprise-level dedicated hosting solution.

For each host, we uploaded and downloaded several large files to measure bandwidth, we called their technical support at various hours to test response time, helpfulness, and how fast they got an answer back to us, and, of course, we rated the prices they charge for the services offered.

At the end of the tests, each host was initially given points out of ten for Product (what they offered for the price), Price (how cheap they were), Support (how fast their customer support responded, how much they were willing to support, and how fast we got answers back), and Performance (speed of file transfer and web serving. We

PAUL HUDSON
has been testing out the services of the UK's top web hosts – if you have your own website or plan to have one soon, read on...

then converted these points into a position score out of six to make direct comparison more easy – the scores are from 1st to 6th for each category, where 1st is best. The Overall score at the bottom is also a positional score based upon the sum of the original points totals as well as any miscellaneous factors we encountered (speed of setup, software configuration, etc).

Hanging on the telephone

When calling technical support, we did our best to sound like newbies to Linux, and the questions we asked technical support were: "How do I add a new user?", "I've written a PHP script and it's complaining about an error – can you help?", and "I think my server is running very slow, what's wrong?" In order to simulate the latter two problems, a PHP script was created with just one line, and the one line had an error in – obvious enough for anyone to spot, even if they could program (**phpinfo** was spelt with two Fs), and to make the server a little sluggish we ran the command **yes > /dev/null &** several times, which quickly chews up 100% of the CPU time. Both of these latter two checks are very easy to spot and solve, and it's simply a test as to whether the support people bother to do the very basic checks required to solve the problems.

Please note that all prices we quote here exclude VAT.

HOSTWAY

Hostway has long been quite an odd player in the industry – despite having a large user base and a moderately speedy connectivity backbone, few people I talk to even seem to know it exists. That's a shame, really, because Hostway set up a dedicated server for our tests faster than everyone else – a very quick response time is a distinct advantage. Furthermore, they sent over a very detailed service level agreement (SLA), guaranteeing 99.97%+ networking uptime, which, if not met, puts you in line for an automatic refund.

The server they provided had a 1.8GHz CPU, with 512MB RAM, 2 x 18GB SCSI hard disks, and a 15GB monthly bandwidth limit. Extras included two-tier firewall with

INFORMATION

WEB www.hostway.com/
SETUP £345
MONTHLY £388

Tripwire support, and daily backup of 30GB, which brings the cost to £345 setup, then £388 a month.

The upload and download speeds were amongst the slowest in these tests, with 125KB/s upload and 25KB/s download – quite a disappointment, particularly for a server of this calibre. With regards to technical support, Hostway did fairly well – the maximum time before I was talking to a real person about my "problems" was 50 seconds. Their responses weren't all that hot, though – for example, asking how to add a new user led to this conversation:

Hostway: "I don't myself know too much about Linux, so I'll have to check"

<on hold for 90 seconds>

WHICH HOST?

Hostway: "I've been told it's the useradd function"

LXF: "Oh, okay. How do I use it?"

Hostway: "The guy didn't say"

LXF: "So...?"

Hostway: "I don't know."

Furthermore, the response to the scripting question was "we don't support scripts, but if you want you can send the error message over and our scripting guy might have a look

VERDICT

HOSTWAY

Product 4th

Price 5th

Support 6th

Performance 6th

OVERALL: 6th

at it if he has the time" – a reasonable answer, but a minor annoyance. They won through on the "slow server" issue, sending the correct answer back by email within an hour.

Overall, when you take every factor into consideration, Hostway performed quite poorly – you don't seem to get an awful lot for your money, other than a remarkably good firewall system. The automatic refund clause in the contract is quite impressive though.

1&1

1&1 has shot to the forefront recently by simply

getting its name around the marketplace – whenever we ask people to name a hosting company, "1&1" is generally the answer we get. 1&1 responded very quickly with the server setup, and had things up and running in under four hours. The server itself was a 3GHz P4, with 1GB RAM, a 60GB IDE hard drive, and 20GB of traffic included per month. There was also a backup solution added in, bringing the total cost to £99 setup, and £128.50 per month. Without a doubt the 1&1 service offers the best value – 3GHz for £99 is incredible.

The bandwidth on this machine was among the best in the group: 180KB/s upload and 108KB/s download. These

INFORMATION

WEB www.oneandone.co.uk/

SETUP £99

MONTHLY £128.50

VERDICT

1&1

Product 2nd

Price 3rd

Support 5th

Performance 2nd

OVERALL 2nd

are some outstanding results for this price, and shows that 1&1 must have cleverly designed their connectivity – their 7,000 megabits of connectivity is clearly put to good use!

The technical support was below average though, and also the slowest – the maximum time till I spoke to a person was 305 seconds. Furthermore, the calls were littered with hold music as the technicians seemed to struggle to answer the questions. The simplest request, "How do I add a user", led to this conversation:

1&1: "I'll just find out for you."

<on hold for 100 seconds>

1&1: "Hello? I'm just checking, I'll be with you in just a second..."

<on hold for another 150 seconds>

1&1: "Hello? I'm just going to transfer you to someone now."

<on hold for another 45 seconds>

1&1: "Can I get your account number from you again?"

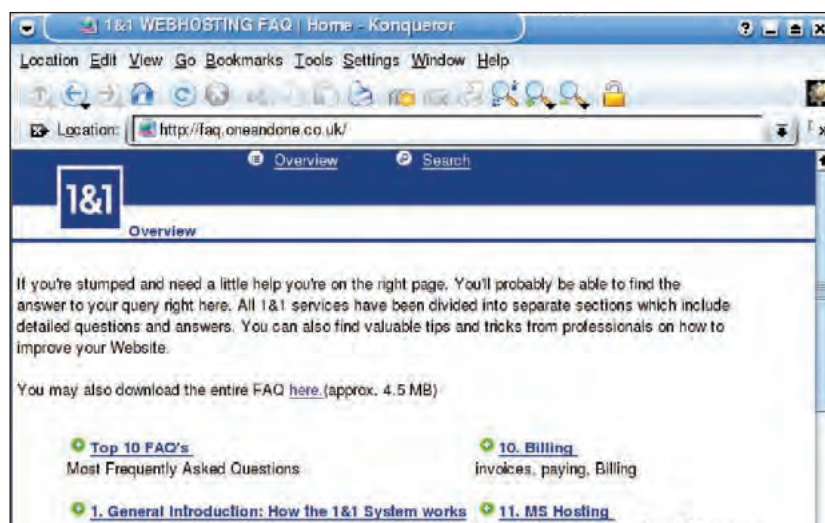
1&1: "OK. Um. I'm just going to put you on hold for one second..."

<on hold for another 35 seconds>

1&1: "We don't offer any technical support for within the terminal."

The last person I spoke to gave me unofficial support, saying it wasn't really covered in the package. However, that was the worst of the conversations by far – 1&1 managed to solve the other two within an hour, sending the response back via email.

Overall, 1&1 offer a value-packed service, but the technical support is almost the worst in the group – perhaps it's a case of "you get what you pay for" for their support lines, which is a shame, because otherwise 1&1 performed excellently.



1&1's FAQ page offers surprisingly detailed advice and solutions for all sorts of problems and questions that you might encounter.

RACKSPACE

EDITOR'S CHOICE

Famed for its "Fanatical Support", Rackspace is a

long-time player in the web hosting industry, hosting some of the most popular sites around – www.theregister.co.uk is just one widely known example. Rackspace took longer than the others to get our server set up, but this was mostly a result of the Fanatical Support team calling me up personally to ask a few questions to make sure I was getting the most for my money.

At the end of the day the server came with a dual 2000+ Athlons, 1GB of RAM, 2 x 18GB SCSI hard drives, 30GB/month bandwidth – no firewall or backup solution, though. One interesting point is that the Rackspace solution

INFORMATION

WEB www.rackspace.co.uk/

SETUP £499

MONTHLY £349

comes with Red Hat Enterprise Linux ES – quite a pricey OS, and certainly a big piece of added value. In total, the Rackspace machine was £499 setup plus £349 a month – it might seem like a high price compared to some of the others here, but considering that this is the most powerful machine in the test with the most powerful OS, backed by the promise of the most comprehensive technical support, and that figure starts to look a lot more justified.

The bandwidth on the machine was a little above average, averaging 152.8KB/s upload and 73.4KB/s download – not the best by a long shot, but still good enough for this level of usage.

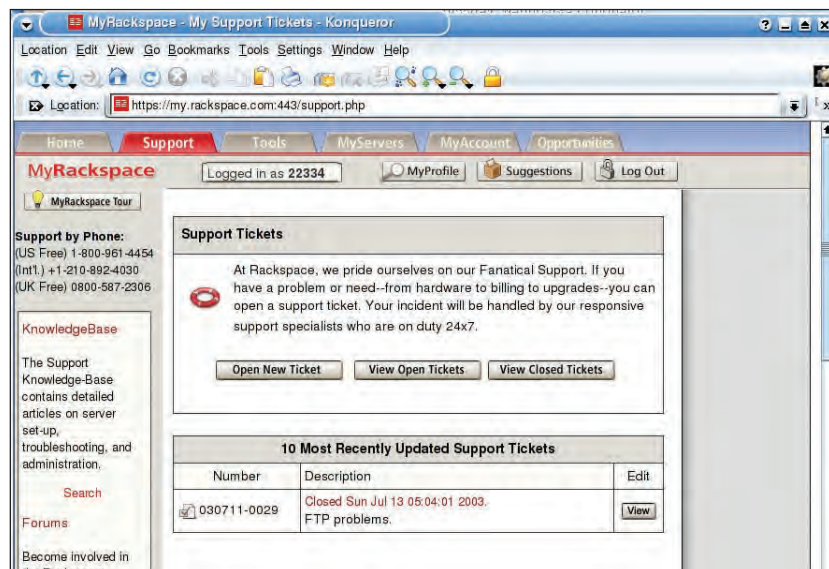
WHICH HOST?

On the technical support front, Rackspace did particularly well – unsurprisingly! The maximum time it took before I was talking to a tech support person was 0 seconds – they answered immediately, with no hold music or “Press 1 for sales, 3 for technical support, etc”. The first time I called up they asked for me for a secret question and answer to verify me on following calls – something none of the other companies did. The technician answered my user-adding question immediately, providing detailed steps for me to follow.

The other calls weren't so good, though – asking about the simple scripting problem resulted in “Sorry we don't support scripts for free, but we can agree a fee to help you if you like”. However, there was worse to come – when calling up and asking about the slow server, I purposefully answered the secret question incorrectly, and the technical support person I was speaking to said “Are you sure you don't mean...”, then *gave me the correct answer*. Since then, Rackspace has implemented a new security policy to make their secret question/answer system be more secure.

To get the answer to the question I was passed over to Rackspace US tech support (although the call remained free for me), who, after putting me on hold for four minutes, found the problem (the multiple “yes” binaries), and instructed me on how to kill them off myself.

All in all, Rackspace was very impressive – the product it offered was far and away the best, and their support was



Rackspace's customer portal allows you to check server status, bandwidth usage, and configuration, as well as file support calls and view existing solutions.

VERDICT RACKSPACE

Product 1st
Price 4th
Support 2nd
Performance 3rd
OVERALL 1st

also very good. The price you pay, while being more than a lot of providers here, is the least expensive high-end solution here, and perhaps the highlight of the package is Rackspace's Fanatical Support – it really is as good as they say. No wonder it has so many high-profile customers!

DESIGNER SERVERS

Designer Servers was one of the two hosting

companies in this test to provide us with a shared server account rather than a dedicated server, leaving us with less control over the server. Easyspace on the opposite page is the other company offering an arrangement of this type. However, we were still able to get full SSH access to do as we pleased with our part of the server. The machine was powered by a dual 800MHz PIII with 1GB of RAM. The relatively modest specification, and the fact that it is after all a shared server, means that Designer Servers price this service starting at just £300 a year – ours was £330 thanks to the addition of a spam and virus filter.

INFORMATION

WEB www.dsrv.co.uk
SETUP n/a
MONTHLY £300/year

VERDICT DESIGNER SERVERS

Product 6th
Price 1st
Support 4th
Performance 1st
OVERALL 4th

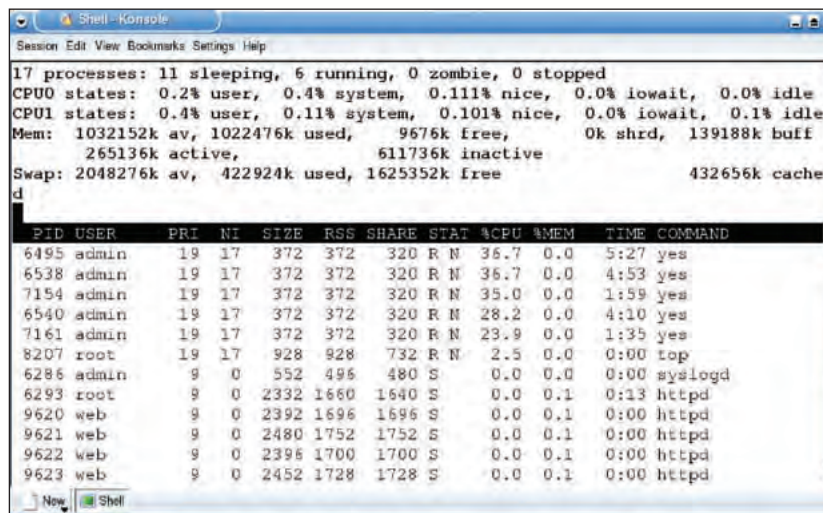
With regards to bandwidth, the machine performed excellent – 148.1KB/s uploading and 149.7KB/s download, which are far and away the best combined results in the test, and Designer Servers deserve particular praise for this.

Designer Servers also did well on technical support, with only a maximum of fifteen seconds spent before speaking to a tech support person. I was hand-held through adding a user, explaining what each command does as well as pointing me towards a website with more information. The response to the broken script problem was: “We don't support third-party scripts, but let me just check it's nothing obvious... oh, yes, there's the problem – you've spelt **phpinfo()** wrong”.

The only fly in the tech support ointment was that the technical support person failed to spot what was chewing up all the CPU time, despite the problem being glaringly obvious – see the screenshot on the left.

In conclusion, Designer Servers' product was the worst in the group – a shared server, and one that leaves customers in a position where they don't really have complete control over it. However, this isn't too much of a problem when you take cost into consideration. Designer Servers' price (five times cheaper than its nearest rival) is a huge difference that could attract self-employed people or very small companies. The high point of Designer Servers is its connectivity – it can really only be summed up as “wow”.

Designer Servers said that there was nothing wrong with the server performance, despite “yes” chewing up so much CPU time.



TDM

TDM are another company which, despite running a large chunk of the UK's websites, don't seem to have much of a public image – again, this is a big shame because their offering packed quite a punch.

The server that TDM provided was a 2.4GHz Xeon with 512MB RAM, 2x36GB SCSI hard drives, and Red Hat 8 pre-installed – all very high-specification equipment, which makes the £299 setup fee and £357 monthly fee a bit of a bargain. Also included in that price is full backup and firewalling, which makes it even more competitive – many rivals (including Designer Servers as mentioned earlier) charge for these extra features. Our tests showed 138KB/s upload and 39.9KB/s download for bandwidth, which is a little below average in the group under examination in this review. The 5GB bandwidth limit isn't terribly high either compared to the others, although it may well be enough for most customer's needs at this level.

INFORMATION

WEB www.tdmgroup.com/
SETUP £299
MONTHLY £357

VERDICT

TDM

Product 3rd

Price 4th

Support 1st

Performance 5th

OVERALL 3rd

The one place where TDM shine through is their technical support – like Rackspace, TDM have no hold music on their tech support line, which meant that I was connected to a person in seconds. Each person I spoke to in all three questions took the problem in their stride, answering in no time at all – they took me through adding a user and setting a password step-by-step, fixed the broken PHP script and told me the problem, and spotted the out of control “yes” processes and killed them for me. This reflects excellently on TDM, and shows they've obviously invested a lot in their staff – it's a shame they don't publicise this!

So, while the bandwidth is about average in the tests, the product is certainly above average, and the tech support is the best – it's hard to beat such a comprehensive and quick tech support response such as TDM gives.

EASYSPEACE

Easyspace took a unique route with its pitch, by

providing us with space on a shared server, as with Designer Servers, but running the whole thing through User Mode Linux, giving us full control over our corner of the machine. Therefore we had 128MB RAM assigned to us, although the rest of the machine was masked behind UML. Included in the deal was full backup and firewalling, with the box running a RAID solution for hot failover.

One particularly interesting feature of the Easyspace package is the lack of a hard bandwidth cap – for £107 a month (no setup fee), you have a 10Mbit line shared with 8 other people, and you can transfer as much as you are able. The soft limit is 500GB – it's hard for the line to take much more given the contention ratio of 8:1. Either way, this is an unrivalled offer by Easyspace, and is a definite boon for people who know they need a lot of data transfer.

The performance was quite different from the others, returning 67.7KB/s for upload and 121KB/s for download –

INFORMATION

WEB www.easyspace.co.uk/
SETUP n/a
MONTHLY £107

VERDICT

EASYSPEACE

Product 5th

Price 2nd

Support 3rd

Performance 4th

OVERALL 5th

very respectable, and among the best in the group. With unmetered data transfer and 121KB/s, even popular sites should have room to breathe. The package that Easyspace gave us was geared towards very heavy users and/or resellers, which means it gets a lot of horsepower behind it.

On the tech support front, Easyspace handled all three of our questions fairly well – however, they were answered by email, which is a minor pain compared to the quick technical support hotlines the other companies provide. What happens if a customer's problem is stopping your access to their mail? One interesting point is that Easyspace made it quite clear it'd prefer me not to use the console – their web management system handles adding users amongst perhaps a hundred other things.

Overall, Easyspace lost out by offering a shared server solution, but won through on the price, and also performed above average for technical support *despite* the fact that tech support is handled by email.

Conclusion

As you can see, pretty much everyone scored highly

in at least one category and poorly in others – the scores are very mixed. This isn't to say that any of the companies were bad in any part of the tests – it's just a comparative rating.

TDM scored very highly on customer support thanks to a strong commitment, and technicians who seemed to enjoy helping out – the speed of TDM's response was incredibly quick, and they smoothly answered all questions put to them. Combined with the fact that the hardware they offer is of a very good standard, they only need to sort out the connectivity – once that's done, they've got it made.

Coming in second place is 1&1, with an excellent offering combined with a price so low you're probably thinking I've misplaced a decimal point. The big sticking point was its

technical support – 1&1 has the longest waiting times, and also didn't managed to solve all the queries. However, even this drawback isn't enough to counter the amount of value you're getting from the service – if you want a great package at for less than £150, 1&1 should be your first stop.

Finally, in first place – quite convincingly – is Rackspace. Not only does it offer the best hardware, but they do so at a very reasonable price – at £349, it's the most economical of the high-end offerings. When it comes to support, Rackspace's commitment to customer satisfaction is almost universally agreed upon, and we're happy to concur – zero waiting times at any time of the day, and you get to speak directly to people who clearly have a massive amount of training and experience in their area. ■